

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Architecture for Distributed Computing System and
Automated Design, Deployment, and Management of
Distributed Applications**

Inventors:

Tilda Brown

Charlie D. Chase

Kevin Grealish

Galen C. Hunt

Aamer Hydrie

David Noble

Geoff Outhred

Glenn Peterson

Bassam Tabbara

Alexander Torone

Robert V. Welland

1 **RELATED APPLICATIONS**

2 This patent application claims the benefit of U.S. Provisional Application
3 No. 60/452,736, filed March 6, 2003, the disclosure of which is incorporated
4 herein by reference.

5 This patent application is also related to the following US patent
6 applications (all of which are incorporated herein by reference):

7 US Patent Application Serial No. 10/382,942, filed on March 6, 2003, titled
8 “Virtual Network Topology Generation”;

9 US Patent Application Serial No. 09/695,812, filed on October 24, 2000,
10 titled “System and Method for Distributed Management of Shared Computers”;

11 US Patent Application Serial No. 09/695,813, filed on October 24, 2000,
12 titled “System and Method for Logical Modeling of Distributed Computer
13 Systems”;

14 US Patent Application Serial No. 09/695,820, filed on October 24, 2000,
15 titled “System and Method for Restricting Data Transfers and Managing Software
16 Components of Distributed Computers”;

17 US Patent Application Serial No. 09/695,821, filed on October 24, 2000,
18 titled “Using Packet Filters and Network Virtualization to Restrict Network
19 Communications”;

20 US Patent Application Serial No. 09/696,707, filed on October 24, 2000,
21 titled “System and Method for Designing a Logical Model of Distributed
22 Computer System and Deploying Physical Resources According to the Logical
23 Model”; and
24
25

1 US Patent Application Serial No. 09/696,752, filed on October 24, 2000,
2 titled "System and Method Providing Automatic Policy Enforcement in a Multi-
3 Computer Service Application".
4

5 **TECHNICAL FIELD**

6 The invention relates to an architecture for a distributed computing system
7 and automated design, deployment, and management of distributed applications on
8 the distributed computing system.
9

10 **BACKGROUND**

11
12 Internet usage has exploded over the past several years and continues to
13 grow. People have become very comfortable with many services offered on the
14 World Wide Web (or simply "Web"), such as electronic mail, online shopping,
15 gathering news and information, listening to music, viewing video clips, looking
16 for jobs, and so forth. To keep pace with the growing demand for Internet-based
17 services, there has been tremendous growth in the computer systems dedicated to
18 hosting Websites, providing backend services for those sites, and storing data
19 associated with the sites.
20

21 One type of distributed computer system is an Internet data center (IDC),
22 which is a specifically designed complex that houses many computers for hosting
23 Internet-based services. IDCs, which also go by the names "Webfarms" and
24 "server farms", typically house hundreds to thousands of computers in climate-
25

1 controlled, physically secure buildings. These computers are interconnected to run
2 one or more programs supporting one or more Internet services or Websites. IDCs
3 provide reliable Internet access, reliable power supplies, and a secure operating
4 environment.

5
6 Fig. 1 shows an Internet data center 100. It has many server computers 102
7 arranged in a specially constructed room. The computers are general-purpose
8 computers, typically configured as servers. An Internet data center may be
9 constructed to house a single site for a single entity (e.g., a data center for Yahoo!
10 or MSN), or to accommodate multiple sites for multiple entities (e.g., an Exodus
11 center that host sites for multiple companies).

12 The IDC 100 is illustrated with three entities—entity A, entity B, and entity
13 C—that share the computer resources. These entities represent various companies
14 that want a presence on the Web. The IDC 100 has a pool of additional computers
15 104 that may be used by the entities at times of heavy traffic. For example, an
16 entity engaged in online retailing may experience significantly more demand
17 during the Christmas season. The additional computers give the IDC flexibility to
18 meet this demand.

19 Today, large IDCs are complex and often called upon to host multiple
20 applications. For instance, some websites may operate several thousand
21 computers, and host many distributed applications. These distributed applications
22 often have complex networking requirements that require operators to physically
23 connect computers to certain network switches, as well as manually arrange the
24 wiring configurations within the IDC to support the complex applications. As a
25 result, this task of building physical network topologies to conform to the

1 application requirements can be a cumbersome, time consuming process that is
2 prone to human error. Accordingly, there is a need for improved techniques for
3 designing and deploying distributed applications onto the physical computing
4 system.

5 6 **SUMMARY**

7 An architecture and methodology for designing, deploying, and managing a
8 distributed application onto a distributed computing system is described.

9 10 **BRIEF DESCRIPTION OF THE DRAWINGS**

11 Similar reference numbers are used throughout the figures to reference like
12 components and/or features.

13 Fig. 1 illustrates an example of an Internet data center.

14 Fig. 2 illustrates an example of a service.

15 Figs. 3-8 illustrate example layer abstractions.

16 Figs. 9-10 illustrate an example SDM type space.

17 Figs. 11-15 illustrate example layer abstractions.

18 Fig. 16 illustrates an example process.

19 Figs. 17-19 illustrate example components as discussed herein.

20 Figs. 20-21 illustrate an example graphical user interface.

21 Fig. 22 illustrates an example SDM model.

22 Fig. 23 illustrates an example deployment.

23 Fig. 24 illustrates example types.

24 Fig. 25 illustrates example instance requests.

25 Fig. 26 illustrates example revalidation of constraints.

1 Fig. 27 illustrates an example logical architecture of an SDM runtime.
2 Fig. 28 illustrates an example graphical representation of a service.
3 Fig. 29 illustrates an example instance space.
4 Fig. 30 illustrates an example of packaging data into an SDU.
5 Fig. 31 illustrates an example type space, member space, and instance
6 space.
7 Fig. 32 illustrates an example member tree.
8 Fig. 33 illustrates an example instance tree.
9 Fig. 34 illustrates an example implementation of the systems described
10 herein.
11 Fig. 35 illustrates example of tracking creation of component instances.
12 Figs. 36-39 illustrate example component instance events.
13 Fig. 40 illustrates an example of a partitioned runtime.
14 Fig. 41 illustrates an example member space.
15 Fig. 42 illustrates an example instance hierarchy.
16 Fig. 43 illustrates an example of partitioning an instance space.
17 Fig. 44 illustrates example relationships between various components.
18 Fig. 45 illustrates an example fixed identity trust relationship.
19 Figs. 46-47 illustrate an example arrangement of components.
20 Fig. 48 illustrates an example platform architecture.
21 Fig. 49 illustrates example usage flow for application deployment.
22 Fig. 50 illustrates example application settings and host settings.
23 Fig. 51 illustrates example phases for a deployment tool.
24 Fig. 52 illustrates an example visualization of a data center description.
25 Figs. 53-54 illustrate example flow diagrams.

1 Fig. 55 illustrates an example of handling an SDU.

2 Figs. 56-58 illustrate example flow diagrams.

3 Fig. 59 illustrates an example model architecture.

4 Fig. 60 illustrates example layers of management.

5 Fig. 61 illustrates an example operation of a system.

6 Fig. 62 illustrates an example connector arrangement.

7 Figs. 63-67 illustrate an example physical configuration of devices.

8 Fig. 68 illustrates an example request graph.

9 Fig. 69 illustrates an example reply graph.

10 Figs. 70-86 illustrates example scenarios in which the invention may be
11 used.

12 Fig. 87 illustrates an example services platform architecture.

13 Fig. 88 illustrates example components in a system.

14 Fig. 89 illustrates example products that may be included in a system
15 described herein.

16 Fig. 90 illustrates various resource management components.

17 Fig. 91 illustrates an example arrangement of multiple LANs.

18 Fig. 92 illustrates an example ADS architecture.

19 Fig. 93 illustrates an example ADS remote boot and imaging system.

20 Fig. 94 illustrates an example topology arrangement.

21 Fig. 95 illustrates an SDML example.

22 Fig. 96 illustrates an example collection of data in a SDU.

23 Fig. 97 illustrates an example of dynamic binding using SDM runtime
24 APIs.

25 Fig. 98 illustrates an example SDM arrangement.

1 Fig. 99 illustrates an example deployment.
2 Fig. 100 illustrates an example system architecture.
3 Fig. 101 illustrates an example of various deployment layers.
4 Fig. 102 illustrates example operations logic.
5 Figs. 103-105 illustrate example changes due to the Internet.
6 Fig. 106 illustrates an example application lifecycle.
7 Fig. 107 illustrates example benefits of a new architecture.
8 Fig. 108 illustrates an example of converting complex systems into simple
9 diagrams.
10 Fig. 109 illustrates an example service.
11 Fig. 110 illustrates an example SQL cluster.
12 Fig. 111 illustrates an example SDM data center model.
13 Fig. 112 illustrates an example design application surface.
14 Fig. 113 illustrates an example SDM service in a data center.
15 Fig. 114 illustrates example resource managers.
16 Fig. 115 illustrates an example of resource virtualization.
17 Fig. 116 illustrates example programming operations logic.
18 Fig. 117 illustrates example interaction with operations logic.
19 Figs. 118-119 illustrate an example of managing heterogeneous
20 environments.

21 22 **DETAILED DESCRIPTION**

23 The following disclosure describes a number of aspects pertaining to an
24 architecture for designing and implementing a distributed computing system with
25 large-scale application services. The disclosure includes discussion of a service

1 definition model (SDM) and an SDM runtime environment. The disclosure
2 further includes design aspects such as how to model data center components, how
3 to model a distributed application description, and techniques for logically placing
4 a modeled application onto a modeled data center and validating this logical
5 placement at design time. The disclosure further explains deployment aspects
6 such as how to instantiate the model using physical resources, physical placement
7 of the distributed application on the physical resources to facilitate application
8 deployment at the physical data center. The disclosure also addresses management
9 aspects, including using the SDM to provide contextual management feedback,
10 tracking, and operations feedback. The disclosure discusses various resource
11 managers used in deployment of the application across physical resources and to
12 support the management aspects.

13 **Service Definition Model (SDM)**

14 The service definition model (SDM) provides tools and a context for an
15 application architect to design distributed computer applications and data centers
16 in an abstract manner. The model defines a set of elements that represent
17 functional units of the applications that will eventually be implemented by
18 physical computer resources and software. Associated with the model elements is
19 a schema that dictates how functional operations represented by the components
20 are to be specified.

22 **SDM Overview**

24 **Introduction**

1
2 Internet Era

3 Over the last decade we have witnessed the Internet emerge as a computing
4 platform. More and more software companies are adopting the “software as a
5 service” model. These services are typically comprised of several components
6 running on many machines including servers, networking equipment and other
7 specialized hardware. Loosely coupled, asynchronous programming models are
8 becoming the norm. Scalability, availability and reliability are critical to the
9 success of these distributed services.
10

11 We are also witnessing a change in hardware trends. High density servers
12 and specialized network hardware are widespread in data centers. Switched fabrics
13 are replacing system buses and providing greater flexibility in system
14 configurations. Hardware cost now plays a small role in the Total Cost of
15 Ownership (TCO) metric compared to the cost of training and maintaining a
16 dedicated operations staff. While rock-solid operational practices are vital to any
17 highly available service, these practices are difficult to repeat consistently because
18 of the fallibility that results from people executing manual procedures.
19

20 In the emerging software as a service era, the focus of development is shifting
21 away from the desktop and toward the server. Along with this change of focus
22 comes a plethora of new problems for software developers, hardware vendors, and
23 IT professionals:
24
25

- 1 ■ Services are larger and more complex – services are time-consuming
2 to develop, difficult and costly to maintain, and risky to extend with
3 additional functionality.
- 4 ■ Services tend to be monolithic – services tend to rely on custom
5 components and specific configurations. Portions of many services
6 cannot be removed, upgraded independently, or replaced with
7 alternatives without impacting the availability of the service.
- 8 ■ Services rely on specific hardware configurations – whether it's a
9 certain network topology or a dependency on a specific network
10 appliance, the binding between hardware and software significantly
11 reduces the ability to host services in different data center
12 environments.
- 13 ■ Services demand operational consistency – most services require a
14 staff of operations personnel to function. The lack of a common
15 platform reduces the ability to reuse code and enact operational best
16 practices across services. Unfortunately, operations staff must be
17 trained in the specifics of each service and retrained as each service
18 evolves.

19 The terms “service” and “application” are used interchangeably throughout
20 this document. In general, an application could be viewed as a collection of
21
22
23
24
25

1 distributed services. For example, Hotmail would be an application comprised of
2 multiple services where each service performs a different function.

3 These problems are not unlike those of the desktop and DOS era (circa
4 1980's). DOS defined valuable core services for application developers such as
5 disk management, file system, console facilities, etc. It did, however, leave many
6 complex tasks up to the ISVs. As an example, WordPerfect and Lotus 123 both
7 independently had to write printer drivers in order to support printing within their
8 respective applications. Similarly, printer hardware vendors had to make deals
9 with the software companies in order to have a successful product. The barrier to
10 entry for ISVs and hardware vendors was exceptionally high. This resulted in only
11 a few successful software and hardware companies during this era.

12
13 Microsoft addressed this problem by creating the Windows platform, which
14 dramatically reduced the barrier to entry. Windows defined an abstraction layer for
15 most hardware devices on the PC platform. This relieved the ISVs from having to
16 worry about supporting specific hardware devices. Windows managed all
17 resources within the PC including memory, disk and network. Windows also came
18 with a wealth of additional services that could be utilized by ISVs. This platform
19 sparked enormous growth in the industry. ISVs that targeted the Windows
20 platform were extremely productive. Many new hardware vendors emerged
21 offering cheaper hardware due to the commoditization effect of having a common
22 platform: Windows.
23
24
25

Service Definition Model (SDM)

SDM Fundamentals

The SDM:

- Defines abstractions that make it easier to design distributed applications / services.
- Enables a framework for reuse and automation of operational practices.
- Simplifies deployment and operations of distributed applications and services.

It can be easier to understand what the SDM is by considering that it captures what today is often seen as a complex diagram on the wall near the operators of a service. In these diagrams a box typically represents a running element of the service and the lines connecting the boxes represent communication paths between the service elements. For example, a load balancer connected to some IIS front-end machines which in turn are connected to one or more middle-tier or back-end services.

Another way to think about the SDM is that it is both a meta-model for the behavior of distributed applications / services and a “live” blueprint of a running application / service in its computing environment. The SDM captures the structure of the application in its computing environment, including its allowable

1 software operations, in a declarative and scale-invariant manner. The ability to
2 declaratively describe the topology of a service, including the bindings between
3 the hardware and network resources, and the valid operations of its software
4 components, is quite powerful.

5 As an analogy, let's look at Microsoft's Common Object Model (COM).
6 COM standardized how components are packaged, registered, activated,
7 discovered, etc. COM mandates strict rules related to lifetime, memory
8 management and interface implementation. These primitives are essential for
9 interoperability because they allow components to be treated as black boxes. COM
10 is the basis for more sophisticated services such as eventing, automation, and
11 OLE.
12

13 Likewise the SDM needs to define some basic primitives on which to build
14 more sophisticated capabilities. These primitives are:

- 15 • Components – units of implementation, deployment and
16 management.
17
- 18 • Ports – named end-points that have an associated type and a set of
19 valid operations.
20
- 21 • Wires – permissible communication paths between ports.
22
- 23 • Layers – separation of resource management ownership and
24 bindings.
25

- Mappings – bindings between components, ports and wires at each layer.

The remainder of this document will describe each of these primitives in more detail.

Components, Ports and Wires

For the purposes of this document, it is useful to consider a graphical representation of a simple service called MyService drawn using components, ports and wires. See Fig. 2. In the diagram, boxes represent components, diamonds represent ports, and dashed lines represent wires.

- MyService is a compound component because it uses the components MyFrontEnd and MyBackEnd.
- MyService has one visible port called web which is a delegated port implemented by the MyFrontEnd component.
- MyFrontEnd has two ports, the delegated port and a port labeled catalog.
- MyBackEnd has one port labeled data.
- The MyFrontEnd and MyBackEnd components have a potential communication relationship that binds the catalog port to the data port through a wire.

Components

Components are units of implementation, deployment and management. Examples of components are a dedicated server running Windows Server, an IIS virtual web site or a SQL database. Components generally have machine boundaries, but are not required to as evidenced by web services hosted on a single IIS server.

Components expose functionality through ports and communicate through wires. Simple components can only have ports as members. Components that use other components are referred to as compound components, which can have ports and wires as members in addition to other components.

Compound components are created through composition and do not have any implementation associated with them. Compound component ports are delegated ports from inner components. Compound components make collocation, encapsulation and reuse possible and so can be thought of as a way to organize an application / service and its behaviors.

Only the public ports of a component are visible outside the component. Compound components to the outside world look like simple components with the internal structure of the components they use hidden through encapsulation. In fact, a simple component could be replaced with a compound component or vice versa as long as the port types and behaviors supported by both are exactly the same.

Ports

Ports are named end-points that define a set of behaviors. Ports have an associated type or role and are typically associated with a set of allowed

1 operations. Examples of ports are an HTTP server port, a SOAP port with a set of
2 allowed operations, etc. Ports can be delegated which means that an outer
3 component can expose the port of an inner component as its own.

4 Ports form the public interface (behavior) to a component. Ports are the only
5 members of a component that can be made public (visible).

6 Wires

7 Wires are permissible bindings between ports and represent topological
8 relationships between ports (and components). Wires do not specify any instance
9 interconnect topology but instead express a “potentiality” for an instance
10 interconnect topology.

11 Wires are essentially busses and can contain one or more port members.
12 Wires should not be mistaken for a point-to-point relationship. A given port cannot
13 appear more than once within the same wire.

14 Schema

15
16 In order to describe an application / service, it is necessary to have a
17 standard schema for the SDM. The SDM schema should be expressible using XSD
18 and XML grammars. While it is beyond the scope of this document to describe the
19 SDM schema in great detail, it is necessary to provide some brief explanation as
20 context for the subject matter described later in this document. Below is a
21 simplified view of the SDM schema.

```
22     <sdm>  
23         <identityReference />  
24         <portClasses />  
25         <wireClasses />  
26         <componentClasses />  
27         <hostRelations />  
28         <portTypes />
```

```
1      <wireTypes />
2      <componentTypes />
3    </sdm>
```

Please read the SDM Schema specification and review the sample XSD files at <http://big/> for more detailed information on the SDM schema.

SDM Class

Every component, port and wire in an application / service is a type created through use of a class. New types can be created from existing classes and types. An SDM class is essentially an abstraction for common features. For example, Web Service can be modeled as class as can a SQL Database. In the MyService application, MyFrontEnd would be a new type derived from the class Web Service; and MyBackEnd would be a new type derived from the class SQL Database.

Below is an example of the class schemas for ports, wires and components.

```
<portClass name="ServerDataAccess" layer="Application">
  <settingSchema>
    <xs:element name="databaseName" type="xs:string" />
  </settingSchema>
</portClass>
```

```
<wireClass name="DataConnection" layer="Application">
  <settingSchema>
    <xs:element name="useSSL" type="xs:boolean" />
  </settingSchema>
```

```
1      <portClassesAllowed>
2          <portClassRef name="ServerDataAccess" maxOccurs="1" />
3          <portClassRef name="ClientDataAccess" />
4      </portClassesAllowed>
5  </wireClass>
```

```
11 <componentClass name="Database" layer="Application">
12     <deploymentSchema>
13         <xs:element name="sqlScriptFilePath" type="xs:string"
14 maxOccurs="unbounded" />
15     </deploymentSchema>
16     <settingSchema>
17         <xs:element name="databaseName" type="xs:string"/>
18     </settingSchema>
19     <portClassesAllowed closed="true">
20         <portClassRef name="ServerDataAccess" />
21     </portClassesAllowed>
22 </componentClass>
```

1 Notice that each componentClass and wireClass schema can contain a setting
2 schema, deployment schema and port classes allowed. The portClass does not
3 have a port classes allowed section. These schemas are defined as follows:

- 4 • Setting Schema is the XSD for the configuration parameters on
5 components, ports and wires that can be design-time validated.
- 6 • Deployment Schema is the XSD that expresses what installation
7 parameters need to be set in order for the component, port or wire to
8 be installed. This manifest could be the schema for Fusion or some
9 other installer technology.
- 10 • Port Classes Allowed is where components and wires declare the
11 allowable ports by referencing declared port classes.

12
13 Please refer to the SDM Schema Design Specification at <http://big> for more details
14 on the class schemas.

15 Class Relationships

16
17 A component, port or wire that can host other components is declared using
18 a hostRelations schema that identifies the installer and the component classes it
19 can host. One can think of the hostRelations element as a directional link between
20 classes where one of the components, ports or wires is acting as a host for the
21 others.

22 Hosting a component means providing the execution environment for a
23 component's code. For example, SQL can be a host for components of class
24 Database as shown in the example below.

25 `<hostRelations>`


```
1      <installer name="DatabaseInstaller" codeType="InstallerPlugIn" />
2      <hostRelation                                classRef="database"
3      componentHostClassRef="host:SQL"  installerRef="DatabaseInstaller"
4      />
5      </hostRelations>
```

6 SDM Types

7
8 There are three distinct spaces that the SDM models: resource, application and
9 instance. The instance space is discussed later in this document. The resource
10 space is where classes live and are the building blocks from which applications are
11 constructed. The application space is where types reside. Below is an example of
12 the XML for port, wire and component types.

```
13      <portType name="UserDataServer" class="ServerDataAccess">
14          <deployment />
15          <settings />
16      </portType>
```

```
17  
18  
19      <wireType name="UserData" class="DataConnection">
20          <deployment />
21          <settings>
22              <useSSL>false</useSSL>
23          </settings>
24          <portTypeRefs>
25              <portTypeRef name="UserDataServer"/>
```

1 <portTypeRef name="UserDataClient"/>
2 </portTypeRefs>
3 </wireType>

13 <componentType name="SQLBackEnd" class="Database">
14 <deployment>
15 <sqlScriptFilePath>%install%\mydatabaseDfn.sql</sqlScriptFilePath>
16 </deployment>
17 <settings>
18 <databaseName>UserData</databaseName>
19 </settings>
20 <ports>
21 <port name="userData" type="UserDataServer"/>
22 </ports>
23 </componentType>

1 Notice each portType, wireType and componentType in the SDM schema contains
2 setting and deployment values.

- 3 • Settings is XML for the settings schema that supplies the
4 configuration values for components, ports and wires and can be
5 design-time validated.
- 6 • Deployment is the XML for the deployment manifest that expresses
7 the values that the configuration parameters need to be set to in order
8 for the component, port or wire to install properly.

9
10 Please refer to the SDM Schema Design Specification at <http://big> for more details
11 on types.

12 Compound Components

13 Compound components can be used to define an application and its topological
14 relationships to other components, ports and wires. Compound components do not
15 have an associated implementation and instead use delegation of ports and host
16 relationships to expose behavior of member components and ports.

17 The XML below shows how the compound component MyService might be
18 described using the SDM.

```
19 <compoundComponentType name="MyService">  
20   <components>  
21     <component name="MyFrontEnd" type="IISFrontEnd" />  
22     <component name="MyBackEnd" type="SQLBackEnd" />  
23   </components>  
24   <wires>  
25
```

```

1      <wire name="data" type="UserData">
2          <members>
3              <member componentName="MyFrontEnd"
4                  portName="serverData" />
5              <member componentName="MyBackEnd"
6                  portName="userData" />
7          </members>
8      </wire>
9  </wires>
10 </compoundComponentType>

```

Instances

While components, ports and wires define the structure and behavior of an application / service, they do not define the running instances. Every component, port and wire type declaration can have one or more instances. Instances are the result of deploying an application / service such that physical resources (servers, network switch ports and disks) are allocated, and software resources (operating systems, runtime hosts, application code) are installed and configured.

It is the job of the SDM Runtime to track all instances from the time of creation until they have been deleted.

SDM Runtime

The SDM Runtime does not itself create instances of components, ports and wires; instead, it provides a set of APIs that are used to coordinate the creation and management of SDM instances. The actual creation of an instance, such as a

1 server running Windows Server with IIS as the host for a web service component,
2 will typically involve multiple entities and could possibly take hours or days to
3 complete.

4 The SDM Runtime knows when a “create SDM instance” process starts and
5 when it terminates, either with success or failure. The SDM Runtime also knows
6 what changes are made to an SDM instance during its lifetime. One way to think
7 about the SDM Runtime is that it is an accountant that records all transactions
8 related to a given application / service SDM such that it can be queried for
9 information about the instances associated with the specified SDM.

10 The first step in creating an SDM instance is registration of an application / service
11 SDM with the SDM Runtime. Once the SDM Runtime knows about a given SDM,
12 the instance creation process can be invoked using Factories and Resource
13 Managers (explained below).

14 Please read the SDM Runtime Architecture specification at <http://big/> for
15 more detailed information on the APIs and runtime design.

16 Hosts and Factories

17 Components that are capable of “hosting” other components are called
18 hosts and act as factories for the classes they support. A component can be
19 declared a host of one or more component classes using the SDM schema
20 hostRelations element described previously.

21 While hosts provide the execution environment for a component’s code,
22 factories are the actual services that create SDM instances of a given type and
23 interact with the SDM Runtime through the SDM Runtime APIs. Factories can
24
25

1 support one or more component classes and must register with the SDM Runtime
2 specifying which component classes they support.

3 It is possible for a given factory to support multiple hosts of the same type
4 with different configurations as well as for individual factories to exist for each
5 type of host configuration. For example, an IIS Factory can support multiple
6 classes such as Web Service and Web Application. Likewise, the SQL Factory can
7 support different database types such as Database, Partitioned Database and
8 Highly Available Database.

9 Factories do not themselves manage physical resources such as storage,
10 network and servers. Factories interact with physical resources (and their logical
11 equivalents) through Resource Managers.

12 Resource Managers

13 Resource Managers manage the physical and logical resources that are (1)
14 discovered or created as part of a bootstrap process or (2) specified through some
15 declarative XML-based description of the physical environment. Resource
16 managers own all storage, network and server resources and expose a common
17 resource management API to process resource allocation requests and to track
18 ownership of these resources.

19 Examples of resource managers are the NRM (Network Resource
20 Manager), the SRM (Storage Resource Manager), and the PRM (PC Resource
21 Manager). Each of these resource managers is responsible for allocation of a
22 physical port or disk or server and the logical resources they expose such as
23 VLANs, logical disk volumes, file shares, web server, etc. Resource managers are
24
25

1 also responsible for programming the physical devices to effect allocation and de-
2 allocation.

3 In order to program the physical hardware, resource managers interact with
4 the hardware through resource providers that hide the implementation details of
5 the hardware device so that, for example, network switches from multiple vendors
6 can be used interchangeably (given that a provider for the manufacturer's device
7 exists). Like the hardware abstraction layer (HAL) and device driver model in
8 Windows, there is an equivalent abstraction layer for the data center environment
9 that spans servers, network and storage devices.

10 Layers and Mappings

11 While components, ports and wires are powerful abstractions when
12 combined with hosts, factories, resource managers and the SDM runtime, they are
13 not sufficient to deploy and manage a distributed application / service. In order to
14 create and manage the physical instances of these logical abstractions, some
15 additional constructs are needed. Those additional constructs are layers and
16 mappings.

17 Layers

18
19 The need for layers is motivated by the desire to perform design-time
20 validation of deployment requirements of an application / service. Fig. 3 shows the
21 layer abstractions defined by the SDM.

- 22
23 • Application layer describes the distributable components, their
24 deployment requirements and constraints, and their communication
25 relationships in the context of an application / service.

- Host layer describes the configuration and policy settings and constraints for hosts such as IIS, CLR and SQL, among others.
- Virtual Data Center (VDC) layer describes the data center environment settings and constraints from the operating system through the network topology down to the servers, network and storage devices.
- Hardware layer is describes the physical data center environment and is either discovered or specified in a declarative manner using XML, for example. This layer is not scale-invariant and therefore not modeled in the SDM, but is included for completeness.

Mappings

Because the SDM is layered, there needs to be a way to bind between the various layers. A mapping is essentially a binding of a component or port at one layer to a component or port at the next layer down. A mapping can be described as follows:

$$M_T = [T_n \rightarrow T_{n-1}] + [T_{n-1} \rightarrow T_{n-2}] + [T_{n-2} \rightarrow T_{n-3}] [\dots]$$

where M represents a mapping and T represents a component, port or wire and n represents the layer. The arrow symbol represents the direction of the mapping which is always from a higher layer to a lower layer.

For example, in Fig. 4 the component at the application layer named MyFrontEnd is mapped to a component at the host layer called IIS. Likewise the component named MyBackEnd is mapped to the SQL component at the host layer.

Design-time Validation

The binding between a component and its host component at the layer below can surface problems to the developer before the application / service is actually deployed in the live data center. These problems can be due to incompatible types, configuration conflicts, mismatched operations, missing topological relationships, etc. Fig. 5 depicts a settings and constraints checking error between a component and its host with regards to authentication.

In Fig. 6, the attempted mapping depicted in the diagram below would result in an error because there is no potential communication relationship between the IIS and SQL components at the deployment layer.

While the mapping from the MyBackEnd component to the SQL host component could have been a valid binding based on the component and host type compatibility and the lack of configuration conflicts, it is invalid because the MyService SDM defined a topological relationship between MyFrontEnd and MyBackEnd that does not exist at the specified deployment layer.

Settings and Constraints Checking

The ability to map from the application layer to the deployment layer (and so on) is quite powerful because it enables design-time validation of a component's settings against a host's constraints; and it also allows validation of a host's settings against a component's constraints.

Fig. 7 shows a more detailed view of the relationships between components and host at the different layers. Notice that there is a binding between a component

1 at one layer and a host component at the next layer down all the way through to
2 the VDC layer.

3 In Fig. 7, MyFrontEnd is a Web Service hosted by IIS which is in turn a
4 Windows Application hosted by Windows Server. There is an IIS factory that
5 supports creation and deletion of Web Service and Web Application component
6 instances just as there is a Windows Application factory that is responsible for
7 creating and deleting instances of IIS and SQL.

8 Fig. 8 shows how design-time validation would work between components
9 at the different layers using the SDM settings and constraints semantics described
10 previously.

11 Notice that the constraints of a component at the layer above are validated
12 against the settings of a host component at the layer below. Also notice that the
13 constraints of the host component are validated against the settings of the
14 component to be hosted.

15 This two-way settings and constraint checking allows a developer to
16 reliably develop his/her application / service in the context of the operational
17 environment described using SDM semantics all the way down. In order to
18 describe a data center such that its description can be relied upon during the
19 development process, it is necessary to create an abstraction of the data center
20 referred to as the VDC.

21 Virtual Data Center (VDC)

22 A VDC is a logical representation of a physical data center environment
23 that simplifies the developer's view of the data center. Ideally an IT Professional
24 or Architect should be able to describe the data center in the same scale-invariant
25

1 manner that a developer can describe a distributed application / service. The way
2 to think about the VDC is that it is an abstraction of the server, network and
3 storage resources within the data center and their topological relationships.
4 A typical data center diagram is quite complex with multiple interconnected
5 servers, network equipment, IP addresses, VLANs, operating systems, storage, etc.
6 all expressed on a single diagram drawn using Visio or a similar tool. In addition
7 to the diagram, there are usually long documents that prescribe exactly how the
8 data center is partitioned, configured and managed.

9 An example of this complexity is the Microsoft Systems Architecture
10 (MSA) Enterprise Data Center (EDC). It should be obvious that keeping the
11 manually drawn diagrams and documents current with the state of the data center
12 over time as updates and upgrades are applied becomes a costly if not impossible
13 task. Likewise, the ability to validate the environment against the document
14 prescriptions is difficult and prone to human error.

15 The ability to represent a complex data center such as the MSA EDC in a
16 scale-invariant manner would be immensely powerful to both the developer and
17 the IT professional. The ability to describe a data center using components, ports
18 and wires provides a powerful framework within which to model and validate
19 deployment requirements that is missing in today's design and deployment
20 process.

21 22 SDM Fundamentals

23 The SDM:
24
25

- Defines abstractions that make it easier to design distributed applications / services.
- Enables a framework for reuse and automation of operational practices.
- Simplifies deployment and operations of distributed applications and services.

It can be easier to understand what the SDM is by considering that it captures what today is often seen as a complex diagram on the wall near the operators of a service. In these diagrams a box typically represents a running element of the service and the lines connecting the boxes represent communication paths between the service elements. For example, a load balancer connected to some IIS front-end machines which in turn are connected to one or more middle-tier or back-end services.

Another way to think about the SDM is that it is both a meta-model for the behavior of distributed applications / services and a “live” blueprint of a running application / service in its computing environment. The SDM captures the structure of the application in its computing environment, including its allowable software operations, in a declarative and scale-invariant manner. The ability to declaratively describe the topology of a service, including the bindings between the hardware and network resources, and the valid operations of its software components, is quite powerful.

As an analogy, let’s look at Microsoft’s Common Object Model (COM). COM standardized how components are packaged, registered, activated,

1 discovered, etc. COM mandates strict rules related to lifetime, memory
2 management and interface implementation. These primitives are essential for
3 interoperability because they allow components to be treated as black boxes. COM
4 is the basis for more sophisticated services such as eventing, automation, and
5 OLE.

6 Likewise the SDM needs to define some basic primitives on which to build
7 more sophisticated capabilities. These primitives are:

- 8 • Components – units of implementation, deployment and
9 management.
- 10 • Ports – named end-points that have an associated type and a set of
11 valid operations.
- 12 • Wires – permissible communication paths between ports.
- 13 • Layers – separation of resource management ownership and
14 bindings.
- 15 • Mappings – bindings between components, ports and wires at each
16 layer.
- 17
- 18

19 The remainder of this document will describe each of these primitives in more
20 detail.

21 Components, Ports and Wires

22
23 For the purposes of this document, it is useful to consider a graphical
24 representation of a simple service called MyService drawn using components,
25 ports and wires.

1 In Fig. 2, boxes represent components, diamonds represent ports, and dashed lines
2 represent wires.

- 3 • MyService is a compound component because it uses the
4 components MyFrontEnd and MyBackEnd.
- 5 • MyService has one visible port called web which is a delegated port
6 implemented by the MyFrontEnd component.
- 7 • MyFrontEnd has two ports, the delegated port and a port labeled
8 catalog.
- 9 • MyBackEnd has one port labeled data.
- 10 • The MyFrontEnd and MyBackEnd components have a potential
11 communication relationship that binds the catalog port to the data
12 port through a wire.
- 13
- 14

15 Components

16
17 Components are units of implementation, deployment and management.
18 Examples of components are a dedicated server running Windows Server, an IIS
19 virtual web site or a SQL database. Components generally have machine
20 boundaries, but are not required to as evidenced by multiple IIS virtual web sites
21 hosted on a single server.

22 Components expose functionality through ports and communicate through
23 wires. Simple components can only have ports as members. Components that use
24 other components are referred to as compound components, which can have ports
25 and wires as members in addition to other components.

1 Compound components are created through composition and do not have
2 any implementation associated with them. Compound component ports are
3 delegated ports from inner components. Compound components make collocation,
4 encapsulation and reuse possible and so can be thought of as a way to organize an
5 application / service and its behaviors.

6 Only the public ports of a component are visible outside the component.
7 Compound components to the outside world look like simple components with the
8 internal structure of the components they use hidden through encapsulation. In
9 fact, a simple component could be replaced with a compound component or vice
10 versa as long as the port types and behaviors supported by both are exactly the
11 same.

12 Ports

13 Ports are named end-points that have an associated type and are typically
14 associated with a protocol role and a set of allowed operations. Examples of ports
15 are an HTTP server port, a SOAP port with a set of allowed operations, etc. Ports
16 can be delegated which means that an outer component can expose the port of an
17 inner component as its own.

18 Ports form the public interface (behavior) to a component. Ports are the
19 only members of a component that can be made public (visible).

20 Wires

21
22 Wires are permissible bindings between ports and represent topological
23 relationships between ports (and components). Wires do not specify any instance
24 interconnect topology but instead express a “potentiality” for an instance
25 interconnect topology.

1 Wires are essentially busses and can contain one or more port members.
2 Wires should not be mistaken for a point-to-point relationship. A given port cannot
3 appear more than once within the same wire.

4 Schema

5 In order to describe an application / service, it is necessary to have a
6 standard schema for the SDM. The SDM schema should be expressible using XSD
7 and XML grammars. While it is beyond the scope of this document to describe the
8 SDM schema in great detail, it is necessary to provide some brief explanation as
9 context for the subject matter described later in this document. Below is a
10 simplified view of the SDM schema.

```
12 <sdm>  
13   <import />  
14   <identityReference />  
15   <information />  
16   <portImplementationType />  
17   <wireImplementationType />  
18   <componentImplementationType />  
19   <hostRelations />  
20   <portTypes />  
21   <wireTypes />  
22   <componentTypes />  
23 </sdm>
```

24 Please read the SDM Schema specification and review the sample XSD files at
25 <http://big/> for more detailed information on the SDM schema.

26 Types

27 Every component, port and wire used in an application / service is a type.
28 Type is essentially equivalent to class in object-oriented languages like C++ and
29 C#, and like it is with classes, new types can be created from existing types. The

1 scale-invariant space is represented in the SDM schema by portTypes, wireTypes
2 and componentTypes. Scale-invariance implies that a component, port or wire can
3 be represented once in an application / service SDM even though there may be
4 multiple instances of each in the actual data center.

5 A type is ultimately derived from an implementation type, which is
6 essentially an abstraction for common technology features. For example, Web
7 Service can be modeled as an implementation type as can SQL Database. In the
8 MyService application, MyFrontEnd would be a new type derived from the
9 implementation type Web Service and MyBackEnd would be a new type derived
10 from the implementation type SQL Database.

11 Each componentImplementationType and wireImplementationType SDM
12 schema element can contain a settings schema, deployment manifest and port
13 implementation reference. The portImplementationType element does not have a
14 port implementation reference. Fig. 9 illustrates what the SDM implementation
15 type space looks like.

- 16 • Settings Schema is the XSD for the configuration parameters on
17 components, ports and wires that can be design-time validated.
- 18 • Deployment Manifest is the XSD that expresses what installation
19 parameters need to be set in order for the component, port or wire to
20 be installed. This manifest could be the schema for Fusion or some
21 other installer technology.
- 22 • Port Implementation Reference is where components and wires
23 declare the allowable ports by referencing declared port
24 implementation types.
- 25

1 In addition, a component that can host other components is declared using a
2 hostRelations SDM schema element that identifies the installer and the component
3 implementation types it can host. One can think of the hostRelations element as a
4 directional link between component implementation types where one of the
5 components is acting as a host for the other component(s). Hosting a component
6 means providing the execution environment for a component's code. For example,
7 IIS is a host for components of implementation type Web Service and Web
8 Application. Hosts will be explained in more detail later in this document.

9 Each portType, wireType and componentType element in the SDM schema
10 contains application constraint values, deployment values and host constraint
11 values. In addition, the wireType element contains a port types element that
12 defines the allowable port types on the specified wire type; and the
13 componentType element contains a hosted types list element that identifies those
14 implementation types that can be hosted on the specified component type. Fig. 10
15 shows the SDM type space.

- 16 • Settings Values is XML for the settings schema that supplies the
17 configuration values for components, ports and wires and can be
18 design-time validated against a host's constraints values.
- 19 • Deployment Values is the XML for the deployment manifest that
20 expresses the values that the configuration parameters need to be set
21 to in order for the component, port or wire to function properly.
- 22 • Constraints Values is the XML for the settings schema that supplies
23 the configuration parameter values that a component, port or wire of
24

1 a host must be set to. Constraints values can be design-time
2 validated against the settings values of the underlying host.

- 3 • Port Types is the XML that lists the allowable port types that can be
4 a member of the specified wire.
- 5 • Hosted Type List is the XML where a component declares the list of
6 component implementation types it can host.
7

8 Instances

9
10 While components, ports and wires define the structure and behavior of an
11 application / service, they do not define the running instances. Every component,
12 port and wire type declaration can have one or more instances. Instances are the
13 result of deploying an application / service such that physical resources (servers,
14 network switch ports and disks) are allocated, and software resources (operating
15 systems, runtime hosts, application code) are installed and configured.

16 It is the job of the SDM Runtime to track all instances from the time of
17 creation until they have been deleted.

18 SDM Runtime

19 The SDM Runtime does not itself create instances of components, ports and
20 wires; instead, it provides a set of APIs that are used to coordinate the creation and
21 management of SDM instances. The actual creation of an instance, such as a
22 server running Windows Server with IIS as the host for a web service component,
23 will typically involve multiple entities and could possibly take hours or days to
24 complete.
25

1 The SDM Runtime knows when a “create SDM instance” process starts and
2 when it terminates, either with success or failure. The SDM Runtime also knows
3 what changes are made to an SDM instance during its lifetime. One way to think
4 about the SDM Runtime is that it is an accountant that records all transactions
5 related to a given application / service SDM such that it can be queried for
6 information about the instances associated with the specified SDM.

7 The first step in creating an SDM instance is registration of an application /
8 service SDM with the SDM Runtime. Once the SDM Runtime knows about a
9 given SDM, the instance creation process can be invoked using Factories and
10 Resource Managers (explained below).

11 Please read the SDM Runtime Architecture specification at <http://big/> for
12 more detailed information on the APIs and runtime design.

13 Hosts and Factories

14 Components that are capable of “hosting” other components are called
15 hosts and act as factories for the implementation types they support. A component
16 can be declared a host of one or more component implementation types using the
17 SDM schema hostRelations element described previously.

18 While hosts provide the execution environment for a component’s code,
19 factories are the actual services that create SDM instances of a given type and
20 interact with the SDM Runtime through the SDM Runtime APIs. Factories can
21 support one or more component implementation types and must register with the
22 SDM Runtime specifying which component implementation types they support.
23 It is possible for a given factory to support multiple hosts of the same type with
24 different configurations as well as for individual factories to exist for each type of
25

1 host configuration. For example, an IIS Factory can support multiple
2 implementation types such as Web Service and Web Application. Likewise, the
3 SQL Factory can support different database types such as Database, Partitioned
4 Database and Highly Available Database.

5 Factories do not themselves manage physical resources such as storage,
6 network and servers. Factories interact with physical resources (and their logical
7 equivalents) through Resource Managers.

8 Resource Managers

9 Resource Managers manage the physical and logical resources that are (1)
10 discovered or created as part of a bootstrap process or (2) specified through some
11 declarative XML-based description of the physical environment. Resource
12 managers own all storage, network and server resources and expose a common
13 resource management API to process resource allocation requests and to track
14 ownership of these resources.

15 Examples of resource managers are the NRM (Network Resource
16 Manager), the SRM (Storage Resource Manager), and the PRM (PC Resource
17 Manager). Each of these resource managers is responsible for allocation of a
18 physical port or disk or server and the logical resources they expose such as
19 VLANs, logical disk volumes, file shares, web server, etc. Resource managers are
20 also responsible for programming the physical devices to effect allocation and de-
21 allocation.

22 In order to program the physical hardware, resource managers interact with
23 the hardware through resource providers that hide the implementation details of
24 the hardware device so that, for example, network switches from multiple vendors
25

1 can be used interchangeably (given that a provider for the manufacturer's device
2 exists). Like the hardware abstraction layer (HAL) and device driver model in
3 Windows, there is an equivalent abstraction layer for the data center environment
4 that spans servers, network and storage devices.

5 Layers and Mappings

6 While components, ports and wires are powerful abstractions when
7 combined with hosts, factories, resource managers and the SDM runtime, they are
8 not sufficient to deploy and manage a distributed application / service. In order to
9 create and manage the physical instances of these logical abstractions, some
10 additional constructs are needed. Those additional constructs are layers and
11 mappings.

12 Layers

13 The need for layers is motivated by the desire to perform design-time
14 validation of deployment requirements of an application / service. Fig. 11 shows
15 the layer abstractions defined by the SDM.
16

- 17
- 18 • Application layer describes the distributable components, their
19 deployment requirements and constraints, and their communication
20 relationships in the context of an application / service.
- 21 • Deployment layer describes the configuration and policy settings
22 and constraints for hosts such as IIS, CLR and SQL, among others.
- 23
- 24 • Virtual Data Center (VDC) layer describes the data center
25 environment settings and constraints from the operating system

1 through the network topology down to the servers, network and
2 storage devices.

- 3 • Hardware layer is describes the physical data center environment
4 and is either discovered or specified in a declarative manner using
5 XML, for example. This layer is not scale-invariant and therefore
6 not modeled in the SDM, but is included for completeness.
7

8 Mappings

9 Because the SDM is layered, there needs to be a way to bind between the
10 various layers. A mapping is essentially a binding of a component or port at one
11 layer to a component or port at the next layer down. A mapping can be described
12 as follows:

$$13 \quad M_T = [T_n \rightarrow T_{n-1}] + [T_{n-1} \rightarrow T_{n-2}] + [T_{n-2} \rightarrow T_{n-3}] [\dots]$$

14 where M represents a mapping and T represents a component, port
15 or wire and n represents the layer. The arrow symbol represents the
16 direction of the mapping which is always from a higher layer to a
17 lower layer.

18 For example, in Fig. 12 the component at the application layer named
19 MyFrontEnd is mapped to a component at the deployment layer called IIS.
20 Likewise the component named MyBackEnd is mapped to the SQL component at
21 the deployment layer.

22 Design-time Validation

23
24 The binding between a component and its host component at the layer
25 below can surface problems to the developer before the application / service is

1 actually deployed in the live data center. These problems can be due to
2 incompatible types, configuration conflicts, mismatched operations, missing
3 topological relationships, etc. For example, the attempted mapping depicted in Fig.
4 13 would result in an error because there is no potential communication
5 relationship between the IIS and SQL components at the deployment layer.

6 While the mapping from the MyBackEnd component to the SQL host
7 component could have been a valid binding based on the component and host type
8 compatibility and the lack of configuration conflicts, it is invalid because the
9 MyService SDM defined a topological relationship between MyFrontEnd and
10 MyBackEnd that does not exist at the specified deployment layer.

11 Settings and Constraints Checking

12 The ability to map from the application layer to the deployment layer (and
13 so on) is quite powerful because it enables design-time validation of a
14 component's settings against a host's constraints; and it also allows validation of a
15 host's settings against a component's constraints.

16 Fig. 14 shows a more detailed view of the relationships between
17 components and host at the different layers. Notice that there is a binding between
18 a component at one layer and a host component at the next layer down all the way
19 through to the VDC layer.

20 In Fig. 14, MyFrontEnd is a Web Service hosted by IIS which is in turn a
21 Windows Application hosted by Windows Server. There is an IIS factory that
22 supports creation and deletion of Web Service and Web Application component
23 instances just as there is a Windows Application factory that is responsible for
24 creating and deleting instances of IIS and SQL.
25

1 Fig. 15 shows how design-time validation would work between
2 components at the different layers using the SDM settings and constraints
3 semantics described previously.

4 Notice that the constraints of a component at the layer above are validated
5 against the settings of the host component at the layer below. Also notice that the
6 constraints of the host component are validated against the settings of the
7 component to be hosted.

8 This two-way settings and constraint checking allows a developer to
9 reliably develop his/her application / service in the context of the operational
10 environment described using SDM semantics all the way down. In order to
11 describe a data center such that its description can be relied upon during the
12 development process, it is necessary to create an abstraction of the data center
13 referred to as the VDC.

14 Virtual Data Center (VDC)

15 A VDC is a logical representation of a physical data center environment
16 that simplifies the developer's view of the data center. Ideally an IT Professional
17 or Architect should be able to describe the data center in the same scale-invariant
18 manner that a developer can describe a distributed application / service. The way
19 to think about the VDC is that it is an abstraction of the server, network and
20 storage resources within the data center and their topological relationships.
21 A typical data center diagram is quite complex with multiple interconnected
22 servers, network equipment, IP addresses, VLANs, operating systems, storage, etc.
23 all expressed on a single diagram drawn using Visio or a similar tool. In addition
24
25

1 to the diagram, there are usually long documents that prescribe exactly how the
2 data center is partitioned, configured and managed.

3 An example of this complexity is the Microsoft Systems Architecture
4 (MSA) Enterprise Data Center (EDC). It should be obvious that keeping the
5 manually drawn diagrams and documents current with the state of the data center
6 over time as updates and upgrades are applied becomes a costly if not impossible
7 task. Likewise, the ability to validate the environment against the document
8 prescriptions is difficult and prone to human error.

9 The ability to represent a complex data center such as the MSA EDC in a
10 scale-invariant manner would be immensely powerful to both the developer and
11 the IT professional. The ability to describe a data center using components, ports
12 and wires provides a powerful framework within which to model and validate
13 deployment requirements that is missing in today's design and deployment
14 process.

15 Agenda: Overview, SDM Building Blocks, Example Application, Example
16 Host, Logical Placement, Deployment, Status.

17 The SDM is a meta-model well-suited for capturing the elemental pieces of
18 distributed applications and their deployment environments. The SDM is
19 authoritative: Application and environment are constructed from their SDM,
20 Changes to the application and environment will be done through the SDM.
21 Provide a namespace for management processes.

22 The Service Definition Model refers to a collection of interrelated schemas:

23 Classes, class relationship and installer schema

24 Component, Port and Wire Types schema

25 Logical placement schema

1 Physical placement schema

2 Instantiation request schema

3 Instance schema

4
5 SDM Classes are the basic building blocks for all distributed applications
6 and deployment environments. Application classes: ASP.Net Web Service,
7 ASP.Net Web Site, BizTalk Orchestration Schedule, Services Components
8 (COM+), etc. Service classes: IIS Server, SQL Server, BizTalk Server. OS,
9 Network & Storage classes: Windows VLAN, Filter, Disk, etc. Hardware classes:
10 Server, Switch, Firewall, Load Balancer, SAN, etc. Classes are authored by
11 system level developers and don't change frequently. Classes are behind every
12 component, port and wire in the SDM. Each class contains a schema for its public
13 settings (simply called settings) and private settings (called deployment).
14 Relationships are captured between classes: component class to port class, wire
15 class to port class, and component class to component class.

16 ASP.Net Web Site Class

```
17 <componentClass name="WebSite" layer="Application">http://big/
18   <settingSchema><xs:schema><xs:complexType><xs:sequence>
19     <xs:element name="webSiteName" type="xs:string"/>
20     <xs:element name="authentication" type="xs:string"/>
21     <xs:element name="sessionState" type="xs:boolean"/>
22   </xs:sequence></xs:complexType></xs:schema></settingSchema>
23   <deploymentSchema><xs:schema><xs:complexType><xs:sequence>
24     <xs:element name="fusionManifest" type="xs:string"/>
25   </xs:sequence></xs:complexType></xs:schema></deploymentSchema>
  <portClassesAllowed closed="true">
    <portClassRef name="ClientDataAccess" />
    <portClassRef name="WebServer" maxOccurs="1"/>
    <portClassRef name="SoapClientInterface" />
    <portClassRef name="RemotingClientInterface" />
  </portClassesAllowed>
</componentClass>
```

SOAP Client Port Class

```
<portClass name="SoapClientInterface" layer="Application">http://big/
  <settingSchema><xs:schema><xs:complexType><xs:sequence>
    <xs:element name="formatter" type="xs:string"/>
    <xs:element name="transport" type="xs:string"/>
  </xs:sequence></xs:complexType></xs:schema></settingSchema>
  <deploymentSchema><xs:schema><xs:complexType><xs:sequence>
    <xs:element name="wsdlFile" type="xs:string"/>
  </xs:sequence></xs:complexType></xs:schema></deploymentSchema>
</portClass>
```

SOAP Wire Class

```
<wireClass name="SoapConnnection" layer="Application">
  <settingSchema/>
  <deploymentSchema/>
  <portClassesAllowed>
    <portClassRef name="SoapServerInterface"/>
    <portClassRef name="SoapClientInterface"/>
  </portClassesAllowed>
</wireClass>
```

IIS Component Class

```
<componentClass name="IIS" layer="Service">
  <settingSchema><xs:schema><xs:complexType><xs:sequence>
    <xs:element name="certificateAuth" type="xs:boolean"/>
    <xs:element name="ntlmAuth" type="xs:boolean"/>
    <xs:element name="sessionStateType" type="xs:string"/>
  </xs:sequence></xs:complexType></xs:schema></settingSchema>
  <deploymentSchema><xs:schema><xs:complexType><xs:sequence>
    <xs:element name="fusionManifest" type="xs:string"/>
  </xs:sequence></xs:complexType></xs:schema></deploymentSchema>
  - <portClassesAllowed>
    <portClassRef name="HTTPServer"/>
    <portClassRef name="HTTPClient"/>
    <portClassRef name="TDSClient"/>
  </portClassesAllowed>
</componentClass>
```

Class Relationships and Installers

```
1      <hostRelation classRef="WebSite" hostClassRef="IIS"
2 installerRef="WebSiteInstaller"/>
3      <installer name="WebSiteInstaller" code="WebSiteInstaller, IISInstaller"
4 codeType="assembly" />
```

5
6 <HostRelation> captures a hosting relationship between classes: IIS can host Web
7 Sites

8 Installers are “plugins” into the SDM Runtime that are responsible for creating a
9 new instances of the component, port and/or wire classes. Installers are also
10 responsible for configuring instances of the class. Different installers might use
11 the same underlying deployment & configuration technology, such as Fusion or
12 WMI.Config.

13 Distributed Application

14 Distributed Applications are constructed from component, port and wire
15 classes. Developers create component, port and wire types from classes. Types
16 are “uses” of classes and supply the values of the setting and deployment schema.
17 Types are a units of reuse. Types map to a single project in Visual Studio.

18 SDM supports composition of types through compound component types.
19 Composition allows bigger distributed applications to be built from smaller ones.
20 Compound component types map to a new project type in Visual Studio –
21 Whitehorse.

22 FMStocks.Web Component Type

```
23      <componentType name="FMStocks.Web" class="WebSite">
24      <ports>
25      <port name="web" type="webServer"/>
      <port name="stock" type="StockClient"/>
      <port name="accounts" type="AccountClient"/>
```

```

1      <port name="trades" type="TradeClient"/>
2    </ports>
3      <settings>
4        <webSiteName>FMStocks.Web</webSiteName>
5        <authentication>Certificate</authentication>
6        <sessionState>true</sessionState>
7      </settings>
8    <deployment>
9      <fusionManifest>fmstocks.web.manifest</fusionManifest>
10    </deployment>
11  </componentType>

```

FMStocks7 Compound Component Type

```

9    <compoundComponentType name="FMStocks">
10      <components>
11        <component name="web" type="FMStocks.Web"/>
12        <component name="svc" type="FMStocks.WebService"/>
13        <component name="biz" type="FMStocks.BusinessService"/>
14        <component name="custdb" type="FMStocks.CustomerDatabase"/>
15      </components>
16      <wires/>
17      <delegatePorts>
18        <port name="web" componentname="web" portname="web"/>
19        <port name="svc" componentname="svc" portname="svc"/>
20      </delegateports>
21    </componentType>

```

SDU and Deployment Environment

Component, port and wire types for a distributed application are packaged along with any binaries in an Service Deployment Unit (SDU). Binaries include all .DLLs, .EXE, .config, static content, etc. SDU represents a portable, independently installable, distributed application. Analogous to the Windows Installer MSI file for Desktop applications. But, unlike desktop applications which primarily target a uniform environment (Windows), distributed applications. Can be hosted on different deployment environments that vary significantly. Must

1 be able to express their requirements on the deployment environment. Must honor
2 all policies of their deployment environment.

3 Therefore, we need a model to express requirements and constraints of both
4 the application and the deployment environment. My WebSite component type
5 needs an IIS server that has been configured with sessions state stored in a SQL
6 database. The web zone will only host webSites components that are using
7 certificate authentication.

9 IIS Component Type

```
10 <componentType name="WebTierIIS" class="IIS">
11   <ports/>
12   <settings>
13     <certificateAuth>true</certificateAuth>
14     <ntlmAuth>>false</ntlmAuth>
15     <sessionStateType>true</sessionStateType>
16   </settings>
17   <deployment/>
18   <hostedClasses>
19     <hostedClass class="WebSite">
20       <!-- constraint language expressed using XPath -->
21       <constraint>/[authentication="certificate"]</constraint>
22     </hostedClass>
23   </hostedClasses>
24 </componentType>
```

19 FMStocks.Web Component Type (revisited)

```
20 <componentType name="FMStocks.Web" class="WebSite">
21   <ports/>
22   <settings>
23     <webSiteName>FMStocks.Web</webSiteName>
24     <authentication>Certificate</authentication>
25     <sessionState>true</sessionState>
26   </settings>
27   <deployment>
28     <fusionManifest>fmstocks.web.manifest</fusionManifest>
29   </deployment>
30   <hostConstraints>
```

```
1      <hostConstraint hostClass="IIS">
2          <constraints>/[sessionStateType="SQL"]</constraints>
3      </hostConstraint>
4  </hostConstraints>
5  </componentType>
```

Logical Placement

Before an SDU can be deployed, we must first do a logical placement of the types on the target deployment environment. Logical placement can be done at design time. Requirements and constraints are checked and the developer is alerted of any errors or warnings. The result of the logical placement is captured in a separate file from the SDU. An SDU can have different logical placements for different deployment environments (Development, Test, Production, etc.) Constraint checking is implemented using XPath and the XSD specified on each component, port and wire class.

Building the Deployment Environment

Deployment environments are built using the SDM model. See Fig. 22. In essence, they are SDM Applications at a different layer. Component, port and wire types are used in the same way to compose service hosts, network architectures, and hardware. In the Whidbey timeframe we will support deploying the application layer only. In ADS V2.0 we will be able to deploy the Service Host, Network and Hardware Layers. Visual studio is building a designer for authoring deployment environments. Visual Studio refers to this as the Logical Infrastructure Model. Fig. 23 illustrates an example deployment.

Instance Request Document

SDM types are scale invariant and can be created to any scale. The Instance Request Document is a declarative definition of the instances that need to

1 be created. Includes the wiring topology. Fig. 24 illustrates example types and
2 Fig. 25 illustrates example instance requests.

3 Physical Placement

4 Physical placement is the act of picking the specific host instance that is the
5 target of deployment. Physical placement is constrained by the logical placement.
6 Constraints are revalidated during physical placement. See Fig. 26.

7 Deployment

8 SDU, Logical Placement file, Instance Request, and Physical Placement file
9 are fed into the SDM Runtime. The SDM Runtime will then invoke the
10 appropriate installer based on the class and host relationship. The installer is
11 responsible for creating a new instance on the host and configuring it to match the
12 settings values on the type. SDM Runtime will maintain a database of all
13 instances, their final setting values, and placement. Runtime API will support
14 querying of the instance space.

16 SDM Schema Design Specification

17 There are three core elements of the SDM schema: ports, wires and
18 components. Ports represent communication endpoints, components represent
19 parts of a distributed application and wires represent communication links between
20 applications parts. These appear in different forms in three separate spaces: the
21 resource space, the application space, and the instance space.

22 In the resource space, the resource classes that applications in the
23 application space will be built from are defined. These classes provide a common
24 categorization of the application parts allowing tool support for a wide range of
25 applications and providing the basis for type checking at design time. We expect

1 these core classes to provide a comprehensive set of features for service design
2 and we expect that they will change slowly over time.

3 In the application space, application types are built. We take a resource
4 class and fill in the details, such as providing links to content, and providing
5 setting values for the properties. We then build distributed applications from these
6 types by associating port types with component types, using component types
7 within a compound component types and describing the communication
8 relationships between the members of a compound component type by using wire
9 types.

10 The instance space consists of the instances created during the process
11 deploying and running an application. We expose the communication relationships
12 we defined in application space through the SDM runtime thus allowing instances
13 to find other instances.

14 Resource Classes

15 We use resource classes to define the elements of application that we need
16 to know about in order to check configuration at design time and then to deploy at
17 run time. These elements are:

- 18 a) Who an application communicates with. In order to verify a distributed
19 application against a network topology we need to know about the
20 protocols that parts of the application can use to communicate with each
21 other. Port classes are used to describe protocol end points. Wire classes are
22 used to describe the relationships that can be constructed between these
23 endpoints.

1 b) What settings apply to an application and how is it deployed.

2 Component classes define building blocks that can be used to construct an
3 application. A component class defines the settings that can be used to
4 control behavior specific to the component and defines a schema for the
5 files and scripts that can be provided to deploy the component.

6 c) What an application depends on in order to function correctly. In order
7 to work correctly a component may depend on certain functionality that
8 must already exist in the target environment. An example is a web service
9 that depends on IIS. We express these requirements as hosting relationships
10 between resources. Using these relationships we can build a dependency
11 tree over the set of resource types that allows us to check ahead of time
12 whether a specific application will run in a particular environment.

13 Application Types

14 We build applications types using the resource classes defined in the
15 resource space. From these classes, we derive port types and wire types to model
16 application specific communication links and we build component types to model
17 the discrete parts of the application.

18 Port types are communications endpoints that describe behavior particular
19 to an application. We take a port resource and provide information that is specific
20 to its use within the application. An example might be a port type that takes a soap
21 resource and provides a WSDL file to define the functions that the application
22 exposes.

23 Wire types define application specific communication paths. A wire type
24 confines a particular wire resource to connecting two compatible application end
25

1 points. For example, we might take a soap wire resource and confine it to
2 connecting the soap port types that we defined above.

3 Component types are used to model the parts of an application that can be
4 deployed independently and can also be distributed across machine boundaries.
5 For example, an application having a web front end and a database backend is
6 likely to consist of several component types. In this case we might take a web
7 service resource and use it to create the web front end component type and a
8 database resource to create the database backend component type. We would then
9 add the appropriate port types to the component types in order to model the
10 application interfaces. We call these port members.

11 Compound component types used are group component types together to
12 form a new component type. A use of a component type inside a compound
13 component is called a component member. We connect the interfaces that
14 component members expose to other members using the wire types we defined
15 earlier. These become the wire members of the compound component.

16 In order for compound components to look like a component, they need to
17 expose interfaces, capabilities and requirements just like a component. We do this
18 by delegating out a subset of the ports members from the component members of
19 the compound component.

20 In order to satisfy the requirements of a component we have to bind that
21 component to another component that has matching capabilities. We call this
22 process binding.

23 Exemplary Implementation

24 In this section we describe the XML schema that we use to define the
25 elements of the SDM model. Settings are used by both applications and resources

1 so we describe them first, then we describe resource classes, then application types
2 and finally the instance space.

3 Naming

4 Namespaces are used to define naming scopes within which classes and
5 types can be defined. Within a namespace all class and type names are unique. A
6 namespace is defined by a name, version, and a cryptographic key that can be used
7 to validate the contents of the namespace.

```
8  
9 <xs:attributeGroup name="identity">  
10   <xs:attribute name="name" type="xs:string" use="required"/>  
11   <xs:attribute name="version" type="fourPartVersionType"  
12     use="required"/>  
13   <xs:attribute name="publicKeyToken" type="publicKeyTokenType"  
14     use="optional"/>  
15 </xs:attributeGroup>
```

16
17 A file version is defined by a four part number of the form N.N.N.N where 0 < N
18 < 65535.

```
19  
20 <xs:simpleType name="fourPartVersionType">  
21   <xs:annotation>  
22     <xs:documentation>Four part version numbers where the segments  
23     are in the range 0-65535 </xs:documentation>  
24   </xs:annotation>  
25   <xs:restriction base="xs:string">  
26     <xs:pattern value="(0|[1-5][0-9]{0,4}|[7-9][0-9]{0,3}|6[0-4][0-  
27     9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-9]?|655|655[0-  
28     2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-9]{0,3}|6[0-4][0-  
29     9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-9]?|655|655[0-  
30     2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-9]{0,3}|6[0-4][0-  
31     9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-9]?|655|655[0-  
32     2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-9]{0,3}|6[0-4][0-  
33     9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-9]?|655|655[0-  
34     2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-9]{0,3}|6[0-4][0-  
35     9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-9]?|655|655[0-  
36     2][0-9]?|655[4-9]|6553[0-5]?)." />
```

```
1 9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-9]?|655|655[0-
2 2][0-9]?|655[4-9]|6553[0-5]?"/>
3   </xs:restriction>
4
5   </xs:simpleType>
```

6 A public key token is a 16 character hex string that identifies the public part of a
7 public/private key pair. The document will be signed using the private key,
8 allowing the user of the document to verify its contents using the public key.

```
9
10 <xs:simpleType name="publicKeyTokenType">
11   <xs:annotation>
12     <xs:documentation>Public Key Token: 16 hex digits in
13     size</xs:documentation>
14   </xs:annotation>
15   <xs:restriction base="xs:string">
16     <xs:pattern value="([0-9][a-f][A-F]){16}"/>
17   </xs:restriction>
18 </xs:simpleType>
```

19 Simple names within the namespace are then constructed using strings.

20 We allow namespaces to reference other namespaces by importing them into the
21 current namespace and then associating an alias with the namespace.

```
22 <xs:complexType name="import">
23   <xs:attribute name="alias" type="xs:string" use="required"/>
24   <xs:attributeGroup ref="identity"/>
25 </xs:complexType>
```

26 References to classes and types are then either simple names that refer to
27 objects defined in the current namespace or compound names that use both an
28 alias and a simple name to identify an object defined in another namespace.

Settings

Both resource classes and application types can expose a settings schema. This schema is used to describe the values that can be provided when a new port, wire or component type is created from a class, when a port type is added to a component type, or when a wire type or component type is used in a compound component type.

Settings schema

We use XSD to describe the settings schema. For the initial release we use a subset of XSD that is limited to simple types and a list of element types.

```
<xs:complexType name="settingSchema">
  <xs:sequence>
    <xs:any namespace="http://www.w3.org/2001/XMLSchema"
      processContents="skip" minOccurs="0"
      maxOccurs="unbounded"/>
  </xs:sequence>
</xs:complexType>
```

Setting Values

Setting values are provided when a type is created based on a class or when a type is used inside a component or compound component. The settings values are a XML block that conforms to the appropriate settings schema.

```
<xs:complexType name="settingValues">
  <xs:sequence>
    <xs:any namespace="##other" processContents="lax"/>
  </xs:sequence>
</xs:complexType>
```

Settings flow

1 We use settings flow to allow settings values to pass from the component
2 type to the members of the component type. Settings flow is implemented using
3 XPATH in setting values sections that select values from the settings schema
4 defined by the type.

5 We identify values that we want to flow contents to by the use of a special
6 attribute that is defined in the SDM namespace. If this attribute exists on an
7 element then we expect the attribute value to be an XPath into the settings schema
8 for the type.

9 Settings Constraints

10 Settings constraints are used to validate and constrain settings values. For
11 example an IIS server may require all web services that it hosts to have some of
12 their settings values confined to a particular value or range of values. We use
13 XPATH to validate settings values (or XQUERY once it is fully supported). We
14 support the following forms of query:

- 15 • Path must exist.
- 16 • Path must not exist.
- 17 • If path exists then [(path must exist | path must not exist)*]

19 Using the first form we can require settings to be set to a particular value or
20 set of values, using the second we can require that a setting not be set to a value or
21 set of values and using the third form we can construct relationships between
22 settings requiring combinations of settings to be set together.

```
23 <xs:complexType name="settingConstraints">  
24   <xs:sequence>  
25     <xs:element name="mustExist" type="simpleTest"  
      minOccurs="0" maxOccurs="unbounded"/>
```



```

1      <xs:element name="mustNotExist" type="simpleTest"
2          minOccurs="0" maxOccurs="unbounded"/>
3      <xs:element name="ifExists" type="nestedTest"
4          minOccurs="0" maxOccurs="unbounded"/>
5      <xs:element name="ifNotExists" type="nestedTest"
6          minOccurs="0" maxOccurs="unbounded"/>
7      </xs:sequence>
8  </xs:complexType>
9
10     <xs:attributeGroup name="testAttributes">
11         <xs:attribute name="path" type="xs:string"/>
12         <xs:attribute name="ifNullPath" type="ifNullPath"/>
13         <xs:attribute name="error" type="xs:int"/>
14         <xs:attribute name="errorDesc" type="xs:string"/>
15     </xs:attributeGroup>
16
17     <xs:complexType name="simpleTest">
18         <xs:attributeGroup ref="testAttributes"/>
19     </xs:complexType>
20
21     <xs:complexType name="nestedTest">
22         <xs:sequence>
23             <xs:element name="mustExist" type="simpleTest"
24                 minOccurs="0" maxOccurs="unbounded"/>
25             <xs:element name="mustNotExist" type="simpleTest"
26                 minOccurs="0" maxOccurs="unbounded"/>
27             <xs:element name="ifExists" type="nestedTest"
28                 minOccurs="0" maxOccurs="unbounded"/>
29             <xs:element name="ifNotExists" type="nestedTest"
30                 minOccurs="0" maxOccurs="unbounded"/>
31         </xs:sequence>
32         <xs:attributeGroup ref="testAttributes"/>
33     </xs:complexType>

```

We need to expose options for dealing with the case when the path does not exist. The following allow the designer to choose to raise an error, insert a value or ignore the test.

```

24     <xs:simpleType name="ifNullPath">
25         <xs:restriction base="xs:string">
26             <xs:enumeration value="skip"/>

```

```
1      <xs:enumeration value="override"/>
2      <xs:enumeration value="returnError"/>
3      </xs:restriction>
4      </xs:simpleType>
```

Examples

The following is a simple schema modeling values a computer class might expose. The schema has a single top level node that identifies the settings group and three properties underneath the node.

```
9      <settingSchema>
10     <xs:schema>
11       <xs:element name="processorSettings">
12         <xs:complexType>
13           <xs:sequence>
14             <xs:element name="numberOfCpus" type="xs:int"/>
15             <xs:element name="memory" type="xs:int" />
16             <xs:element name="dualHomed" type="xs:boolean"/>
17           </xs:sequence>
18         </xs:complexType>
19       </xs:element>
20     </xs:schema>
21   </settingSchema>
```

We could provide the following values for the schema within a type.

```
19     <settings>
20       <processorSettings>
21         <numberOfCpus>4</numberOfCpus>
22         <memory>8000</memory>
23         <dualHomed>>false</dualHomed>
24       </processorSettings>
25     </settings>
```

If we wanted to provide the settings values when the type was used then we would use settings flow.

1 Constraints may be written against these values. In the example, the first is
2 a simple mustExist constraint. The second constraint uses a test to determine
3 whether to evaluate the nested constraints.

```
4  
5  
6  
7  
8  
9 <constraints>  
10   <mustExist path="ProcessorSettings/[memory >= 1000]"  
11     errorDesc="Host machine does not have enough memory"/>  
12  
13   <ifExists path="ProcessorSettings/[cpu >= 2]"  
14     errorDesc="Host machine has two processors but not enough  
15     resources">  
16     <mustExist path="ProcessorSettings/[memory >= 2000]"  
17       errorDesc="Host machine does not have enough memory"/>  
18   </ifExists >  
19 </constraints>
```

20 Resources

21 Base Class

22 All resource class schemas derive from class. They share a settings schema,
23 deployment schema and name and layer attributes. The settings schema describes
24 the settings that apply to types based on this class, the values that they can take
25 and description of each. The deployment schema describes the information that is
required to deploy a type that is based on this resource. The layer attribute

1 associated the resource with one layer in the design space. The name attribute is
2 used to give the class a unique name within the namespace.

```
3  
4 <xs:complexType name="class">  
5   <xs:sequence>  
6     <xs:element name="deploymentSchema" type="deploymentSchema"  
7       minOccurs="0" maxOccurs="1"/>  
8     <xs:element name="settingSchema" type="settingsSchema"  
9       minOccurs="0" maxOccurs="1"/>  
10   </xs:sequence>  
11   <xs:attribute name="name" type="xs:string" use="required"/>  
12   <xs:attribute name="layer" type="layer" use="required"/>  
13 </xs:complexType>
```

14
15 For the deployment schema the namespace is left undefined. The
16 constraints on the schema are entirely the responsibility of the installer for the
17 class.

```
18 <xs:complexType name="deploymentSchema">  
19   <xs:sequence>  
20     <xs:any namespace="##other" processContents="lax"/>  
21   </xs:sequence>  
22 </xs:complexType>
```

23
24 The values provides as part of the deployment section must match the
25 associated deployment schema.

```
26 <xs:complexType name="deploymentValues">  
27   <xs:sequence>  
28     <xs:any namespace="##other" processContents="lax"/>  
29   </xs:sequence>  
30 </xs:complexType>
```

31
32 The layer attribute is an enumeration of four layer types. The application
33 layer contains high level application components such as database and webserver.
34 The service layer contains middleware services such as IIS and SQL. The network
35

1 layer contains operating system, storage and network definitions. The hardware

2 layer contains definitions of the hardware components of a data center.

```
3 <xs:simpleType name="layer">
4   <xs:restriction base="xs:string">
5     <xs:enumeration value="Application"/>
6     <xs:enumeration value="Service"/>
7     <xs:enumeration value="Network"/>
8     <xs:enumeration value="Hardware"/>
9   </xs:restriction>
10 </xs:simpleType>
```

8 Port Class

9
10 Port classes do not contain any information above that defined in the
11 resource base type.

```
12 <xs:complexType name="portClass">
13   <xs:complexContent>
14     <xs:extension base="class">
15     </xs:extension>
16   </xs:complexContent>
17 </xs:complexType>
```

16 Component Class

17 A component class extends the base class by adding a list of allowed port
18 classes.

```
19 <xs:complexType name="componentClass">
20   <xs:complexContent>
21     <xs:extension base="class">
22       <xs:sequence>
23         <xs:element name="portClassesAllowed"
24           type="portClassesAllowed"
25           minOccurs="0" maxOccurs="1"/>
26       </xs:sequence>
27     </xs:extension>
28   </xs:complexContent>
29 </xs:complexType>
```

The list of port classes can be open or closed, if it is closed, then only those port types based on classes that appear in the list can be used on the associated component type. The minOccurs and maxOccurs attributes define the number of times one of these port types can be used.

```
<xs:complexType name="portClassesAllowed">
  <xs:sequence>
    <xs:element name="portClassRef" minOccurs="0"
maxOccurs="unbounded">
      <xs:complexType>
        <xs:attribute name="name" type="xs:string" use="required"/>
        <xs:attribute name="minOccurs" type="xs:int" use="optional"/>
        <xs:attribute name="maxOccurs" type="xs:string" use="optional"/>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
  <xs:attribute name="closed" type="xs:boolean"
    default="true" use="optional"/>
</xs:complexType>
```

Wire Class

The wire class also extends the base schema by adding a list of allowed port classes. In this case the list defines the classes of the port types that may be associated with the wire type.

```
<xs:complexType name="wireClass">
  <xs:complexContent>
    <xs:extension base="class">
      <xs:sequence>
        <xs:element name="portClassesAllowed"
type="portClassesAllowed" minOccurs="0" maxOccurs="1"/>
      </xs:sequence>
    </xs:extension>
  </xs:complexContent>
</xs:complexType>
```

```
1      </xs:extension>
2      </xs:complexContent>
3      </xs:complexType>
```

4 Hosting Relationship

5 A hosting relationship defines is a triple identifying a source class, a target
6 class and an installer. The existence of the relationship indicates that an instance of
7 a type based on the source class could be created using an instance of a type based
8 on the target class and the installer associated with the relationship. The target
9 class must be a component class.

10 For example a webservice class may be the source class in a hosting
11 relationship with an IIS class and the webservice installer. In this case the
12 relationship indicates that it may be possible to create an instance of type
13 MyWebservice on type MyIIS using the installer. We do not know whether it will
14 be possible to create the relationship until we have evaluated constraints that exist
15 in both the application space and the instance space.

```
16 <xs:complexType name="hostRelation">
17   <xs:attribute name="classRef" type="xs:string" use="required"/>
18   <xs:attribute name="componentHostClassRef" type="xs:string"
19   use="required"/>
20   <xs:attribute name="installerRef" type="xs:string" use="required"/>
21 </xs:complexType>
```

22 The installer is identified by name, code type and a link to the binary that
23 implements the installer.

```
24 <xs:complexType name="installer">
25   <xs:sequence>
26     <xs:element name="binary" type="xs:string" minOccurs="1"
27     maxOccurs="1"/>
28   </xs:sequence>
29   <xs:attribute name="codeType" type="xs:string" use="required"/>
```

```
1      <xs:attribute name="name" type="xs:string" use="required"/>
2      </xs:complexType>
```

3 Examples

4 These examples are excerpts from the extended four layer example. See the
5 complete example files for details.

6 First we create some port classes to model access to a database. In this case
7 we have a server port and a client port.

```
8      <portClass name="ServerDataAccess" layer="Application">
9          <settingSchema>
10             <xs:schema>
11                 <xs:complexType>
12                     <xs:sequence>
13                         <xs:element name="databaseName" type="xs:string"/>
14                         <!-- other connection string properties -->
15                     </xs:sequence>
16                 </xs:complexType>
17             </xs:schema>
18          </settingSchema>
19      </portClass>
20      <portClass name="ClientDataAccess" layer="Application"/>
```

16 We then create a wire class that models the communication link between
17 the two port classes. The wire class has some settings and references the two port
18 classes defined above. In this case the wire constrains there to be only one server
19 on the connection, modeling the fact that the client port does not know how to load
20 balance connections across multiple servers. A more complex wire implementation
21 may allow multiple servers and implement some form of management to resolve
22 connections.

```
23
24      <wireClass name="DataConnection" layer="Application">
25          <settingSchema>
```



```

1      <xs:schema>
2          <xs:complexType>
3              <xs:sequence>
4                  <xs:element name="useSSL" type="xs:boolean"/>
5              </xs:sequence>
6          </xs:complexType>
7      </xs:schema>
8  </settingSchema>
9  <portClassesAllowed>
10     <portClassRef name="ServerDataAccess" maxOccurs="1"/>
11     <portClassRef name="ClientDataAccess"/>
12 </portClassesAllowed>
13 </wireClass>

```

Finally we create a component class that models a database. This class has both a settings and a deployment schema and identifies the ports that can exist on a component type based on this class.

```

12 <componentClass name="Database" layer="Application">
13     <deploymentSchema>
14         <xs:schema>
15             <xs:complexType>
16                 <xs:sequence>
17                     <xs:element name="sqlScriptFilePath" type="xs:string"
18                         maxOccurs="unbounded"/>
19                 </xs:sequence>
20             </xs:complexType>
21         </xs:schema>
22     </deploymentSchema>
23     <settingSchema>
24         <xs:schema>
25             <xs:complexType>
26                 <xs:sequence>
27                     <xs:element name="databaseName" type="xs:string"/>
28                 </xs:sequence>
29             </xs:complexType>
30         </xs:schema>
31     </settingSchema>
32     <portClassesAllowed closed="true">
33         <portClassRef name="ServerDataAccess"/>
34     </portClassesAllowed>

```

1 </componentClass>

2 All these components need mapping to compatible host types. In this case
3 SQL server acts as a host for the server port and the database and IIS acts as a host
4 for the sql client port. These classes are defined in a separate namespace aliased to
5 middleware.

6 <hostRelations>
7
8 <installer name="DatabaseInstaller" codeType="InstallerPlugin"/>
9
10 <hostRelation classRef="database"
11 componentHostClassRef="middleware:SQL"
12 installerRef="DatabaseInstaller"/>
13
14 <hostRelation classRef="ServerDataAccess"
15 componentHostClassRef=" middleware:SQL"
16 installerRef="DatabaseInstaller"/>
17
18 <hostRelation classRef="ClientDataAccess"
19 componentHostClassRef=" middleware:IIS"
20 installerRef="WebServiceInstaller"/>
21
22 </hostRelations>

23 Applications

24 The application developer creates component, port and wire types in the
25 application space to model his application. These types are created by selecting
classes that match the layer that the developer is working within and then
supplying values for the classes.

26 Application base type

27 All application type schemas are based on the following application base
28 schema. The base schema attributes identify the class that the type is based on and

1 the name of the type. In the body of the schema we identify the deployment values
2 that will allow this type to be deployed, and the settings for the settings schema on
3 the associated class. The type may define also define a new settings schema that
4 identifies values that can be provided when this type is used within other types.
5 Finally the base type includes a section for host constraints. This section identifies
6 constraints on possible hosts for this type based on the host relationships that exist
7 in the resource space for the class associated with this type.

```
8 <xs:complexType name="baseType">
9   <xs:sequence>
10     <xs:element name="deployment" type="deploymentValues"
11       minOccurs="0" maxOccurs="1"/>
12     <xs:element name="settings" type="settingsValues"
13       minOccurs="0" maxOccurs="1"/>
14     <xs:element name="settingSchema" type="settingSchema"
15       minOccurs="0" maxOccurs="1"/>
16     <xs:element name="hostConstraints" type="hostConstraints"
17       minOccurs="0" maxOccurs="1"/>
18   </xs:sequence>
19   <xs:attribute name="class" type="xs:string" use="required"/>
20   <xs:attribute name="name" type="xs:string" use="required"/>
21 </xs:complexType>
```

22 The hostConstraints section includes a set of constraints for each of the
23 classes that could host the class associated with this type. These classes are
24 identified by the host relations in the resource space. The constraints associated
25 with each class are in terms of the settings schema the classes. The form of the
constraints was defined above.

```
<xs:complexType name="hostConstraints">
```

```

1      <xs:sequence>
2          <xs:element name="hostConstraint" minOccurs="1" maxOccurs="1">
3              <xs:complexType>
4                  <xs:sequence>
5                      <xs:element name="constraint" type="settingConstraint"/>
6                  </xs:sequence>
7                  <xs:attribute name="host" type="xs:string" use="required"/>
8              </xs:complexType>
9          </xs:element>
10         </xs:sequence>
11     </xs:complexType>

```

Port Type

Port types simply use the base type. There is no further information associated with a port type.

```

12     <xs:complexType name="portType">
13         <xs:complexContent>
14             <xs:extension base="baseType">
15                 </xs:extension>
16             </xs:complexContent>
17         </xs:complexType>

```

Wire Type

Wire types extend the base type to add a list of allowed port types. Uses of these port types may then be associated with a use of the wire type within a compound component. By defining wire types in this way, the application designer can constrain the set of allowable connections between parts of his application by only creating wires types for compatible port types.

```

22     <xs:complexType name="wireType">
23         <xs:complexContent>
24             <xs:extension base="baseType">
25                 <xs:sequence>
26                     <xs:element name="portTypeRefs" minOccurs="0">
27                         <xs:complexType>
28                             <xs:sequence>

```

```

1      <xs:element name="portTypeRef"
2          minOccurs="0" maxOccurs="unbounded">
3          <xs:complexType>
4              <xs:attribute name="name" type="xs:string"
5                  use="required"/>
6          </xs:complexType>
7      </xs:element>
8  </xs:sequence>
9  </xs:complexType>
10 </xs:element>
11 </xs:sequence>
12 </xs:extension>
13 </xs:complexContent>
14 </xs:complexType>

```

Component Type

A component type extends the base type to add a list of port members and a list of hosted classes.

Each port member is a use of an existing port type. The list of hosted classes identifies the classes that this component can host. These classes are a subset of the classes identified by the host relationships in the resource space, where this type's class is identified as a potential host.

```

17 <xs:complexType name="componentType">
18   <xs:complexContent>
19     <xs:extension base="baseType">
20       <xs:sequence>
21         <xs:element name="ports" type="portsList"
22             minOccurs="0" maxOccurs="1"/>
23         <xs:element name="hostedClasses" type="hostedClassesList"
24             minOccurs="0" maxOccurs="1"/>
25       </xs:sequence>
26     </xs:extension>
27   </xs:complexContent>
28 </xs:complexType>

```

Each port member in the ports list is identified by name and type. The port name must be unique within the component. The port type must have an associated port class that is allowed on the component class associated with this component type. For each port member we can provide a list of settings that match the schema defined by the port type.

```
<xs:complexType name="portsList">
  <xs:sequence>
    <xs:element name="port" minOccurs="0" maxOccurs="unbounded">
      <xs:complexType>
        <xs:sequence>
          <xs:element name="settings" type="settingValues" minOccurs="0"
maxOccurs="1"/>
        </xs:sequence>
        <xs:attribute name="name" type="xs:string" use="required"/>
        <xs:attribute name="type" type="xs:string"/>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

For each class in the hosted classes list we can associate a list of constraints. These constraints are written with respect to the setting schema of the hosted class.

```
<xs:complexType name="hostedClassesList">
  <xs:sequence>
    <xs:element name="hostedClass" minOccurs="1" maxOccurs="1">
      <xs:complexType>
        <xs:sequence>
          <xs:element name="constraints" type="settingConstraints"
minOccurs="1" maxOccurs="1"/>
        </xs:sequence>
        <xs:attribute name="class" type="xs:string" use="required"/>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
```

Compound Component Type

A compound component type (hereafter referred to as compound component) defines a new component type. When defining the compound component, there is the option to specify that the members of the type should be co-located. If the members are co-located, then when the type is deployed all the members of the type must be deployed on a single host. The compound component also contains a list of component members, a list of wire members, a section defining the ports that the component delegates and a list identifying the classes that the component can host.

```
<xs:complexType name="compoundComponentType">
  <xs:sequence>
    <xs:element name="components" type="components"
      minOccurs="0" maxOccurs="1"/>
    <xs:element name="wires" type="wires"
      minOccurs="0" maxOccurs="1"/>
    <xs:element name="delegatePorts" type="delegatePorts"
      minOccurs="0" maxOccurs="1"/>
    <xs:element name="delegateHostedClasses"
      type="delegateHostedClasses"
      minOccurs="0" maxOccurs="1"/>
  </xs:sequence>
  <xs:attribute name="name" type="xs:string" use="required"/>
  <xs:attribute name="colocate" type="xs:boolean"
    use="optional" default="false"/>
</xs:complexType>
```

The component list identifies uses of component types that have already been defined – we call these the component members of the compound

1 component. Each member has a unique name within the compound component, a
2 reference to the type that defines it and a flag that indicates whether it is singleton
3 or not.

4 If a component member is marked as singleton, then there can only be once
5 instance of this component member within an instance of the containing
6 compound component. If it is not marked as singleton, then instances of a member
7 may be created and deleted according to external factors such as load changes.
8 This means that any component member that is connected to a non-singleton
9 member may see one or more instances of that member at runtime.

10 Each component member may also provide settings values for the settings
11 schema defined in the associated component type.

```
12 <xs:complexType name="components">  
13   <xs:sequence>  
14     <xs:element name="component" minOccurs="0"  
15       maxOccurs="unbounded">  
16       <xs:complexType>  
17         <xs:sequence>  
18           <xs:element name="settings" type="settingValues"  
19             minOccurs="0" maxOccurs="1"/>  
20         </xs:sequence>  
21         <xs:attribute name="name" type="xs:string" use="required"/>  
22         <xs:attribute name="type" type="xs:string" use="required"/>  
23         <xs:attribute name="singleton" type="xs:boolean"  
24           use="optional" default="false"/>  
25       </xs:complexType>  
26     </xs:element>  
27   </xs:sequence>  
28 </xs:complexType>
```

23 A use of wire type within a compound component is called a wire member.
24 Each wire member has a name that is unique to the compound component and
25

1 identifies an associated wire type. Wire member can also provide settings values
2 for the settings schema defined in the wire type.

3 The key role of a wire member is to identify connection between
4 component members within the compound component. The way this is done is to
5 add port references to the wire member. Each port reference identifies a port on a
6 component member within the compound component. The port types of the
7 references ports must match the port types that are associated with the wire type.

```
8 <xs:complexType name="wires">
9   <xs:sequence>
10     <xs:element name="wire" minOccurs="0" maxOccurs="unbounded">
11       <xs:complexType>
12         <xs:sequence>
13           <xs:element name="settings" type="settingValues" minOccurs="0"
14             maxOccurs="1"/>
15           <xs:element name="members" minOccurs="1" maxOccurs="1">
16             <xs:complexType>
17               <xs:sequence>
18                 <xs:element name="member" type="componentPortRef"
19                   minOccurs="0" maxOccurs="unbounded"/>
20               </xs:sequence>
21             </xs:complexType>
22           </xs:element>
23         </xs:sequence>
24       <xs:attribute name="name" type="xs:string" use="required"/>
25       <xs:attribute name="type" type="xs:string"/>
26     </xs:complexType>
27   </xs:element>
28 </xs:sequence>
29 </xs:complexType>
```

1 A port reference identifies a component member within the same containing
2 compound component. The port name is the name of a port member on the
3 component type associated with the component member.

```
4 <xs:complexType name="componentPortRef">  
5   <xs:attribute name="componentName" type="xs:string"/>  
6   <xs:attribute name="portName" type="xs:string" use="required"/>  
7 </xs:complexType>
```

8 A compound component cannot use port types directly as there is no code
9 associated with the compound component that the port member could bind to.
10 Instead we delegate out port members from the component members of the
11 compound component. This means that these ports appear as though they belong
12 to the compound component when it is used as a component type.

13 When a port is delegated, it is identified by first identifying the component
14 member and then the port member within that component. The port can be
15 renamed as part of this process in order to avoid name clashes in cases where ports
16 with the same name are delegated from different component members.

```
17 <xs:complexType name="delegatePorts">  
18   <xs:sequence>  
19     <xs:element name="delegatePort" minOccurs="0"  
20       maxOccurs="unbounded">  
21       <xs:complexType>  
22         <xs:attribute name="name" type="xs:string"/>  
23         <xs:attribute name="componentName" type="xs:string"/>  
24         <xs:attribute name="portName" type="xs:string" use="optional"/>  
25       </xs:complexType>  
26     </xs:element>  
27   </xs:sequence>  
28 </xs:complexType>
```

1 In order to construct hosts that may provide services for a range of different
2 classes we allow a compound component to expose the hosted class declarations
3 from its component members. When the compound component is used as a
4 component type, it then appears that the compound component can act as a host
5 for all the declared classes.

6 To expose these hosted class declarations we use delegation in a similar
7 way to the way in which we delegated port members. We identify the component
8 member that contains the hosted class, and then we identify the class that the
9 component claims to be able to host.

```
10 <xs:complexType name="delegateHostedClasses">  
11   <xs:sequence>  
12     <xs:element name="hostedClassRef"  
13       minOccurs="1" maxOccurs="unbounded">  
14       <xs:complexType>  
15         <xs:attribute name="componentName" type="xs:string"/>  
16         <xs:attribute name="hostedClass" type="xs:string"  
17           use="required"/>  
18       </xs:complexType>  
19     </xs:element>  
20   </xs:sequence>  
21 </xs:complexType>
```

18 Binding

19 Binding is the process where we identify hosts for the members of a
20 particular compound component. We do this in order to check compatibility
21 between an application and the environment in which it will be hosted and to
22 deploy the application. Both the application and the host environment are modeled
23 using compound components so the process of binding is to find matching
24 members from both components that support the connection topology between the
25 members.

1 To identify compatible hosts for a member, we start by looking at the
2 relationship between classes in the resource space. We look at the type of wire or
3 component member and then identify the class associated with the member. We
4 then look for component members in the host component that have compatible
5 classes associated with their component types. We then look at the host constraints
6 on the type associated with the member and see if they match the settings on the
7 host member's type. We then do the reverse, checking the hostedClass constraints
8 on the host member's type against the settings on the type of the member that we
9 want to host.

10 If we are trying to match a component member then we need to check that
11 all the port members of the component member's type can also be hosted on any
12 potential host for the component member.

13 If we are trying to match a wire member, then we have to match any
14 component members that exist on the path between the hosts that we choose for
15 component members in the compound component that we are trying to host.

16 Based on the port classes we described in the previous example, we create
17 two port types.

```
18 <portType name="UserDataServer" class="ServerDataAccess">  
19   <deployment/>  
20   <settings/>  
21 </portType>  
22  
23 <portType name="UserDataClient" class="ServerDataAccess">  
24   <deployment/>  
25   <settings/>  
26 </portType>
```

24 These types are complimented by a wire type.

```
25 <wireType name="UserData" class="DataConnection">
```

```

1      <deployment/>
2      <settings>
3          <useSSL>>false</useSSL>
4      </settings>
5      <portTypeRefs>
6          <portTypeRef name="UserDataServer"/>
7          <portTypeRef name="UserDataClient"/>
8      </portTypeRefs>
9      </wireType>

```

Now we create a component type based on the database class. The database type exposes one server data port.

```

10     <componentType name="UserData" class="Database">
11         <deployment>
12             <sqlScriptFilePath>%install%\mydatabaseDfn.sql</sqlScriptFilePath>
13         </deployment>
14         <settings>
15             <databaseName>UserData</databaseName>
16         </settings>
17         <ports>
18             <port name="userData" type="UserDataServer"/>
19         </ports>
20     </componentType>

```

We could create a compound component type that uses some of these types. The following compound component uses three component types. The first type UserPages represents a web service with two access points, the second type QueryManagement is a middle tier logic component, and the last type is our database type. We connect these components up using two wire types: UserData and QueryManager. The data wire connects the middle tier to the database and the query wire connects the frontend to the middle tier. We then expose two ports: signup and enquiry, from the front end using delegation.

```

24     <compoundComponentType name="UserManagementApplication">
25         <components>

```

```

1      <component name="userPages" type="UserPages"/>
2      <component name="queryLogic" type="QueryManagement"/>
3      <component name="userData" type="UserData" singleton="true"/>
4      </components>
5      <wires>
6          <wire name="data" type="UserData">
7              <members>
8                  <member componentName="queryLogic" portName="userData"/>
9                  <member componentName="userData" portName="userData"/>
10             </members>
11         </wire>
12         <wire name="query" type="QueryManager">
13             <members>
14                 <member componentName="userPages"
15                 portName="queryManager1"/>
16                 <member componentName="userPages"
17                 portName="queryManager2"/>
18                 <member componentName="queryLogic"
19                 portName="queryManager"/>
20             </members>
21         </wire>
22     </wires>
23     <delegatePorts>
24         <delegatePort name="signup" componentName="userPages"
25         portName="signup"/>
26         <delegatePort name="enquiry" componentName="userPages"
27         portName="enquiry"/>
28     </delegatePorts>
29 </compoundComponentType>

```

SDM Document structure

An SDM document has a strong identity which defines the namespace of the document. It imports a list of references other namespaces. The document also contains a information section that identifies document specific attribute such as the document owner, company name and revision date. It then contains lists of port, wire and component classes, followed by a list of host relationships, followed in turn by lists of port, wire and component types.

```

1      <xs:element name="sdm">
2          <xs:annotation>
3              <xs:documentation>SDM root element. It is a container for SDM
4              types.</xs:documentation>
5          </xs:annotation>
6          <xs:complexType>
7              <xs:sequence>
8                  <xs:element name="import" type="import" minOccurs="0"
9                  maxOccurs="unbounded"/>
10                 <xs:element name="information" type="information"
11                 minOccurs="0" maxOccurs="1"/>
12                 <xs:element name="portClasses" minOccurs="0" maxOccurs="1">
13                     <xs:complexType>
14                         <xs:sequence>
15                             <xs:element name="portClass" type="portClass"
16                             minOccurs="1" maxOccurs="unbounded"/>
17                         </xs:sequence>
18                     </xs:complexType>
19                 </xs:element>
20                 <xs:element name="wireClasses" minOccurs="0" maxOccurs="1">
21                     <xs:complexType>
22                         <xs:sequence>
23                             <xs:element name="wireClass" type="wireClass"
24                             minOccurs="1" maxOccurs="unbounded"/>
25                         </xs:sequence>
26                     </xs:complexType>
27                 </xs:element>
28                 <xs:element name="componentClasses" minOccurs="0"
29                 maxOccurs="1">
30                     <xs:complexType>
31                         <xs:sequence>
32                             <xs:element name="componentClass" type="componentClass"
33                             minOccurs="1" maxOccurs="unbounded"/>
34                         </xs:sequence>
35                     </xs:complexType>
36                 </xs:element>
37                 <xs:element name="hostRelations" minOccurs="0"
38                 maxOccurs="1">
39                     <xs:complexType>
40                         <xs:sequence>
41                             <xs:element name="installer" type="installer" minOccurs="1"
42                             maxOccurs="unbounded"/>

```

```

1      <xs:element name="hostRelation" type="hostRelation"
2      minOccurs="1" maxOccurs="unbounded"/>
3      </xs:sequence>
4      </xs:complexType>
5      </xs:element>
6      <xs:element name="portTypes" minOccurs="0" maxOccurs="1">
7      <xs:complexType>
8      <xs:sequence>
9      <xs:element name="portType" type="portType"
10     minOccurs="0" maxOccurs="unbounded"/>
11     </xs:sequence>
12     </xs:complexType>
13     </xs:element>
14     <xs:element name="wireTypes" minOccurs="0" maxOccurs="1">
15     <xs:complexType>
16     <xs:sequence>
17     <xs:element name="wireType" type="wireType"
18     minOccurs="0" maxOccurs="unbounded"/>
19     </xs:sequence>
20     </xs:complexType>
21     </xs:element>
22     <xs:element name="componentTypes" minOccurs="0"
23     maxOccurs="1">
24     <xs:complexType>
25     <xs:sequence>
26     <xs:element name="componentType" type="componentType"
27     minOccurs="0" maxOccurs="unbounded"/>
28     <xs:element name="compoundComponentType"
29     type="compoundComponentType" minOccurs="0"
30     maxOccurs="unbounded"/>
31     </xs:sequence>
32     </xs:complexType>
33     </xs:element>
34     </xs:sequence>
35     <xs:attributeGroup ref="identity"/>
36     </xs:complexType>
37     </xs:element>

```


1 Associated XSD

2 The following is an example structure for a change request.

```
3     <?xml version="1.0" encoding="utf-8" ?>
4     <xs:schema targetNamespace="urn:schemas-microsoft-
5 com:sdmChangeRequest" xmlns="urn:schemas-microsoft-
6 com:sdmChangeRequest" xmlns:settings="urn:schemas-microsoft-
7 com:sdmSettings" xmlns:mstns="http://tempuri.org/XMLSchema.xsd"
8 xmlns:xs="http://www.w3.org/2001/XMLSchema"
9 elementFormDefault="qualified" version="0.7" id="sdmChangeRequest">
10     <xs:import namespace="urn:schemas-microsoft-com:sdmSettings"
11 schemaLocation="SDM7Settings.xsd" />
12     <xs:import namespace="urn:schemas-microsoft-com:sdmNames"
13 schemaLocation="SDM7Names.xsd" />
14     <xs:complexType name="ChangeRequestType">
15         <xs:sequence>
16             <xs:element name="group" type="groupType"
17 minOccurs="0" maxOccurs="unbounded" />
18         </xs:sequence>
19     </xs:complexType>
20     <xs:complexType name="groupType">
21         <xs:sequence>
22             <xs:element name="group" type="groupType"
23 minOccurs="0" maxOccurs="unbounded" />
24             <xs:element name="addInstance"
25 type="addInstanceType" minOccurs="0" maxOccurs="unbounded" />
26             <xs:element name="updateInstance"
27 type="updateInstanceType" minOccurs="0" maxOccurs="unbounded" />
28             <xs:element name="deleteInstance"
29 type="deleteInstanceType" minOccurs="0" maxOccurs="unbounded" />
30             <xs:element name="addConnection"
31 type="addConnectionType" minOccurs="0" maxOccurs="unbounded" />
32             <xs:element name="deleteConnection"
33 type="deleteConnectionType" minOccurs="0" maxOccurs="unbounded" />
34         </xs:sequence>
35         <xs:attribute name="canLeConcurrentlyExecuted"
36 type="xs:boolean" />
37     </xs:complexType>
38     <xs:complexType name="addInstanceType">
39         <xs:sequence>
40             <xs:element name="classSettings"
41 type="settings:settingValues" minOccurs="0" />
42         </xs:sequence>
43     </xs:complexType>
44 </xs:schema>
```

```

1      <xs:element name="typeSettings"
type="settings:settingValues" minOccurs="0" />
2      <!-- setting values for class -->
3      <!-- setting values for type -->
4      </xs:sequence>
5      <xs:attribute name="parent" type="reference" use="optional"
6      />
7      <xs:attribute name="host" type="reference" use="optional"
8      />
9      <xs:attribute name="member" type="xs:string"
10     use="optional" />
11     <xs:attribute name="type" type="xs:string" use="optional" />
12     <xs:attribute name="name" type="xs:string" use="optional"
13     />
14     <!-- the parent of this instance -->
15     <!-- the host of this instance -->
16     <!-- Name of the member on the parent type -->
17     <!-- Fully qualified type that this is an instance of -->
18     <!-- alias for the id that can be filled in when the instance is
created.
19     this name must be unique for all instances of the same
20     member. -->
21     </xs:complexType>
22     <!-- what can we change about an instance? -->
23     <xs:complexType name="updateInstanceType">
24     <xs:sequence>
25     <xs:element name="classSettings"
type="settings:settingValues" minOccurs="0" />
26     <xs:element name="typeSettings"
type="settings:settingValues" minOccurs="0" />
27     <!-- setting values for class -->
28     <!-- setting values for type -->
29     </xs:sequence>
30     <xs:attribute name="id" type="reference" use="required" />
31     <xs:attribute name="parent" type="reference" use="optional"
32     />
33     <xs:attribute name="host" type="reference" use="optional"
34     />
35     <xs:attribute name="member" type="xs:string"
36     use="optional" />
37     <xs:attribute name="type" type="xs:string" use="optional" />
38     <xs:attribute name="name" type="xs:string" use="optional"
39     />

```

```

1      <!-- Unique identifier scoped to the SDM Runtime. This is
generated by t_u101 ? SDM runtime
2          and is immutable -->
3      <!-- the parent of this instance -->
4      <!-- the host of this instance -->
5      <!-- Name of the member on the parent type -->
6      <!-- Fully qualified type that this is an instance of -->
7      <!-- alias for the id that can be filled in when the instance is
created.
8          this name must be unique for all instances of the same
member. -->
9      </xs:complexType>
10     <xs:complexType name="deleteInstanceType">
11         <xs:attribute name="id" type="reference" use="required" />
12         <xs:attribute name="option" type="deleteOptionType"
13 use="required" />
14         <!-- Unique identifier scoped to the SDM Runtime. This is
generated by the SDM runtime
15             and is immutable - _cf2 >
16     </xs:complexType>
17     <xs:complexType name="addConnectionType">
18         <xs:attribute name="port" type="reference" use="required"
19 />
20         <xs:attribute name="wire" type="reference" use="required"
21 />
22     </xs:complexType>
23     <xs:complexType name="deleteConnectionType">
24         <xs:attribute name="port" type="reference" use="required"
25 />
26         <xs:attribute name="wire" type="reference" use="required"
27 />
28     </xs:complexType>
29     <!-- reference can be guid or path -->
30     <xs:simpleType name="reference">
31         <xs:union></xs:union>
32     </xs:simpleType>
33     <!-- delete options are: ??? -->
34     <xs:simpleType name="deleteOptionType">
35         <xs:union></xs:union>
36     </xs:simpleType>
37 </xs:schema>

```

1 The following is an example structure for classes.

```
2     <?xml version="1.0" encoding="utf-8" ?>
3     <xs:schema targetNamespace="urn:schemas-microsoft-com:sdmClasses"
4     xmlns="urn:schemas-microsoft-com:sdmClasses" xmlns:names="urn:schemas-
5     microsoft-com:sdmNames" xmlns:settings="urn:schemas-microsoft-
6     com:sdmSettings" xmlns:xs="http://www.w3.org/2001/XMLSchema"
7     elementFormDefault="qualified" version="0.7" id="sdmClasses">
8         <xs:import namespace="http://www.w3.org/2001/XMLSchema" />
9         <xs:import namespace="urn:schemas-microsoft-com:sdmSettings"
10        schemaLocation="SDM7Settings.xsd" />
11        <xs:import namespace="urn:schemas-microsoft-com:sdmNames"
12        schemaLocation="SDM7Names.xsd" />
13        <!-- TODO [BassamT]: Normalize the port class refs, port type refs
14        and port members on wire classs, wire types and wire members -->
15        <!-- TODO [BassamT]: Is the layer attribute mandatory on a class? -
16        -->
17        <!-- TODO [BassamT]: Add keys and keyefs for validation -->
18        <!-- TODO [BassamT]: Add support for inlined types -->
19        <!-- TODO [BassamT]: scrub minOccurs and maxOccurs -->
20        <!-- TODO [BassamT]: New name for "class", possibly
21        "deployment" -->
22        <!-- TODO [BassamT]: New name for "host", possibly "provider" --
23        >
24        <!-- REVIEW [BassamT]: Can we merge the definitions of port,
25        component, wire classs in this XSD. It would make it less verbose at the cost more
26        semantic analysis. -->
27        <!-- CONSIDER [BassamT]: General attribute mechanism for things
28        like Singleton, Colocation, Inline. -->
29        <!-- TODO [BassamT]: Bindings: member to component member --
30        >
31        <!-- TODO [geoffo]: ports - are they singleton? -->
32        <!-- TODO [geoffo]: delegation - how do we combine ports? -->
33        <!-- TODO [geoffo] Add back <any> in appropriate places -->
34        <!--
35        =====
36        =====
37        = -->
38
39        <!-- SDM root element -->
40        <!--
41        =====
42        =====
43        = -->
44
45        <xs:element name="sdmClasses">
```

```

1      <xs:complexType>
2          <xs:sequence>
3              <xs:element name="import"
4                  type="names:import" minOccurs="0" maxOccurs="unbounded" />
5              <xs:element name="information"
6                  type="information" minOccurs="0" />
7              <xs:element name="portClasses"
8                  minOccurs="0">
9                  <xs:complexType>
10                     <xs:sequence>
11                         <xs:element
12                             name="portClass" type="portClass" maxOccurs="unbounded" />
13                     </xs:sequence>
14                 </xs:complexType>
15             </xs:element>
16             <xs:element name="componentClasses"
17                 minOccurs="0">
18                 <xs:complexType>
19                     <xs:sequence>
20                         <xs:element
21                             name="componentClass" type="componentClass" maxOccurs="unbounded" />
22                     </xs:sequence>
23                 </xs:complexType>
24             </xs:element>
25             <xs:element name="protocols"
26                 minOccurs="0">
27                 <xs:complexType>
28                     <xs:sequence>
29                         <xs:element
30                             name="protocol" type="protocol" maxOccurs="unbounded" />
31                     </xs:sequence>
32                 </xs:complexType>
33             </xs:element>
34             <xs:element name="hostRelations"
35                 minOccurs="0">
36                 <xs:complexType>
37                     <xs:sequence>
38                         <xs:element
39                             name="installer" type="installer" maxOccurs="unbounded" />
40                         <xs:element
41                             name="hostRelation" type="hostRelation" maxOccurs="unbounded" />
42                     </xs:sequence>
43                 </xs:complexType>

```

```

1         </xs:element>
2         </xs:sequence>
3         <xs:attributeGroup ref="names:namespaceIdentity" />
4     </xs:complexType>
5 </xs:element>
6 <!-- SDM type library information -->
7 <xs:complexType name="information">
8     <xs:annotation>
9         <xs:documentation>Human readable information
10        about the SDM type library.</xs:documentation>
11    </xs:annotation>
12    <xs:sequence>
13        <xs:element name="friendlyName" type="xs:string"
14        minOccurs="0" />
15        <xs:element name="companyName" type="xs:string"
16        minOccurs="0" />
17        <xs:element name="copyright" type="xs:string"
18        minOccurs="0" />
19        <xs:element name="trademark" type="xs:string"
20        minOccurs="0" />
21        <xs:element name="description" type="xs:string"
22        minOccurs="0" />
23        <xs:element name="comments" type="xs:string"
24        minOccurs="0" />
25    </xs:sequence>
26 </xs:complexType>
27 <!--
28 =====
29 =====
30 = -->
31
32     <!-- Classes -->
33     <!--
34 =====
35 =====
36 = -->
37
38     <xs:complexType name="baseClass">
39         <xs:sequence>
40             <xs:element name="deploymentSchema"
41             type="settings:deploymentSchema" minOccurs="0" />
42             <xs:element name="settingSchema"
43             type="settings:settingSchema" minOccurs="0" />
44             <!-- XSD schema that for how a class is deployed -->
45             <!-- Setting schema -->

```

```

1      </xs:sequence>
2      <xs:attribute name="name" type="xs:string" use="required"
3  />
4      <xs:attribute name="layer" type="xs:string" use="required"
5  />
6      <!-- REVIEW [BassamT] Are these layers just for benefit of
7  tools, or are they
8  strictly enforced in the SDM model? There are
9  cases where mixing components from different
10 layers makes
11 sense. For example, the filter component might
12 be a
13 a component meta type hosted ISA server which
14 lives in layer 3.
15 However, we want to use the Filter meta-type in
16 layer 2. -->
17 </xs:complexType>
18 <!-- port class -->
19 <xs:complexType name="portClass">
20   <xs:complexContent>
21     <xs:extension base="baseClass" />
22   </xs:complexContent>
23 </xs:complexType>
24 <!-- Component class -->
25 <xs:complexType name="componentClass">
26   <xs:complexContent>
27     <xs:extension base="baseClass">
28       <xs:sequence>
29         <xs:element
30           name="portClassesAllowed" minOccurs="0">
31             <xs:complexType>
32               <xs:sequence>
33                 <xs:element
34                   name="portClassRef" minOccurs="0" maxOccurs="unbounded" />
35               </xs:sequence>
36             </xs:complexType>
37             <xs:attribute
38               name="closed" type="xs:boolean" use="optional" default="true" />
39             <!-- Whether the allowable
40             ports is closed list -->
41             <!-- If this value is "true"
42             then the list of ports is non-extensible. If this value is "false" then the list of ports
43             is open-ended, the ports listed will be considered mandatory. -->
44           </xs:complexType>

```

```

1                                     </xs:element>
2                                     <!-- this will specify a set of constraints
on the set of allowable ports
                                     that can show up on a component
3 type of this meta type. -->
                                     </xs:sequence>
4                                     </xs:extension>
5                                     </xs:complexContent>
6                                     </xs:complexType>
7                                     <xs:complexType name="portClassRef">
                                     <xs:attribute name="name" type="xs:string" use="required"
8 use="required" />
                                     <xs:attribute name="required" type="xs:boolean"
9 use="required" />
                                     <xs:attribute name="singleton" type="xs:boolean"
10 <!-- singleton implies that there can only be one instance of
this port within the parents scope -->
                                     </xs:complexType>
11                                     <!--
12 =====
13 = -->
14                                     <!-- relations -->
15                                     <!--
16 =====
17 = -->
18                                     <xs:complexType name="relation">
19                                     <xs:attribute name="name" type="xs:string" use="required"
20 />
21                                     <xs:attribute name="installer" type="xs:string"
22 use="optional" />
23                                     </xs:complexType>
24                                     <!-- a protocol is a relationship between one or more port classes -->
25                                     <xs:complexType name="protocol">
                                     <xs:complexContent>
                                     <xs:extension base="relation">
                                     <xs:sequence>
                                     <xs:element name="portClassRef"
type="portClassRef" maxOccurs="unbounded" />
                                     </xs:sequence>
                                     </xs:extension>

```



```

1      </xs:complexContent>
2      </xs:complexType>
3      <!-- defines the host relationship between two classes -->
4      <xs:complexType name="hostRelation">
5          <xs:complexContent>
6              <xs:extension base="relation">
7                  <xs:attribute name="classRef" type="xs:string"
8                      use="required" />
9                  <xs:attribute name="hostClassRef"
10                      type="xs:string" use="required" />
11              </xs:extension>
12          </xs:complexContent>
13      </xs:complexType>
14      <!-- the installer type identifies the code responsible for instantiating
15      a relationship -->
16      <xs:complexType name="installer">
17          <xs:sequence>
18              <xs:element name="binary" type="xs:string" />
19          </xs:sequence>
20          <xs:attribute name="codeType" type="xs:string"
21              use="required" />
22          <xs:attribute name="name" type="xs:string" use="required"
23              />
24      </xs:complexType>
25  </xs:schema>

```

1 The following is an example structure for a deployment unit.

```
2     <?xml version="1.0" encoding="UTF-8" ?>
3     <xs:schema targetNamespace="urn:schemas-microsoft-com:sdmSDU"
4     xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:names="urn:schemas-
5     microsoft-com:sdmNames" xmlns="urn:schemas-microsoft-com:sdmSDU"
6     elementFormDefault="qualified" version="0.7" id="sdmSDU">
7         <xs:import namespace="http://www.w3.org/2001/XMLSchema" />
8         <xs:import namespace="urn:schemas-microsoft-com:sdmNames"
9     schemaLocation="SDM7Names.xsd" />
10        <!-- an sdm deployment unit imports one or more sdm type files the
11     includes mappings for a subset of the types from the imported file -->
12        <xs:element name="sdmDeploymentUnit">
13            <xs:annotation>
14                <xs:documentation>
15                    The sdu contains a mapping of SDM types to
16     their implementation.
17                </xs:documentation>
18            </xs:annotation>
19            <xs:complexType>
20                <xs:sequence>
21                    <xs:element name="import"
22     type="names:import" minOccurs="0" maxOccurs="unbounded" />
23                    <xs:element name="implementation"
24     type="implementationMap" minOccurs="0" maxOccurs="unbounded" />
25                </xs:sequence>
26            </xs:complexType>
27        </xs:element>
28        <!-- a description of this deployment unit -->
29        <xs:complexType name="deploymentDescription">
30            <xs:attribute name="name" type="xs:string" />
31            <xs:attribute name="dateCreated" type="xs:string" />
32            <xs:attribute name="creator" type="xs:string" />
33        </xs:complexType>
34        <!-- a mapping from a type to an implementation of the type -->
35        <xs:complexType name="implementationMap">
36            <xs:sequence>
37                <xs:element name="version" type="xs:string"
38     minOccurs="0" maxOccurs="unbounded" />
39            </xs:sequence>
40            <xs:attribute name="type" type="xs:string" />
41            <xs:attribute name="path" type="xs:string" />
42        </xs:complexType>
43    </xs:schema>
```

1 The following is an example structure for instances.

```
2     <?xml version="1.0" encoding="utf-8" ?>
3     <xs:schema targetNamespace="urn:schemas-microsoft-com:sdmInstances"
4     xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:settings="urn:schemas-
5     microsoft-com:sdmSettings" xmlns="urn:schemas-microsoft-com:sdmInstances"
6     elementFormDefault="qualified" version="0.7" id="sdmInstances">
7         <xs:import namespace="http://www.w3.org/2001/XMLSchema" />
8         <xs:import namespace="urn:schemas-microsoft-com:sdmSettings"
9     schemaLocation="SDM7Settings.xsd" />
10        <xs:element name="sdmInstances">
11            <xs:complexType>
12                <xs:sequence>
13                    <xs:element name="import" type="import"
14    minOccurs="0" maxOccurs="unbounded" />
15                    <xs:element name="portInstances"
16    minOccurs="0">
17                        <xs:complexType>
18                            <xs:sequence>
19                                <xs:element
20    name="portInstance" type="portInstance" minOccurs="0"
21    maxOccurs="unbounded" />
22                            </xs:sequence>
23                        </xs:complexType>
24                    </xs:element>
25                    <xs:element name="wireInstances"
26    minOccurs="0">
27                        <xs:complexType>
28                            <xs:sequence>
29                                <xs:element
30    name="wireInstance" type="wireInstance" minOccurs="0"
31    maxOccurs="unbounded" />
32                            </xs:sequence>
33                        </xs:complexType>
34                    </xs:element>
35                    <xs:element name="componentInstances"
36    minOccurs="0">
37                        <xs:complexType>
38                            <xs:sequence>
39                                <xs:element
40    name="componentInstance" type="componentInstance" minOccurs="0"
41    maxOccurs="unbounded" />
42                            </xs:sequence>
43                        </xs:complexType>
44                    </xs:element>
45                </xs:sequence>
46            </xs:complexType>
47        </xs:element>
48    </xs:schema>
```

```

1      name="compoundComponentInstance" type="compoundComponentInstance"
2      minOccurs="0" maxOccurs="unbounded" />
3          </xs:sequence>
4          </xs:complexType>
5          </xs:element>
6      </xs:sequence>
7      </xs:complexType>
8  </xs:element>
9  <xs:complexType name="import">
10     <xs:attribute name="alias" type="xs:string" use="required" />
11     <xs:attributeGroup ref="identity" />
12 </xs:complexType>
13 <!-- ===== Instance Schema =====>
14 -->
15     <xs:complexType name="instanceBase">
16         <xs:sequence>
17             <xs:element name="classSettings"
18 type="settings:settingValues" minOccurs="0" />
19             <xs:element name="typeSettings"
20 type="settings:settingValues" minOccurs="0" />
21             <!-- setting values for class -->
22             <!-- setting values for type -->
23         </xs:sequence>
24         <xs:attribute name="id" type="guid" use="required" />
25         <xs:attribute name="parent" type="guid" use="optional" />
26         <xs:attribute name="host" type="guid" use="optional" />
27         <xs:attribute name="member" type="xs:string"
28 use="optional" />
29         <xs:attribute name="type" type="xs:string" use="required" />
30         <xs:attribute name="name" type="xs:string" use="optional"
31 />
32         <!-- Unique identifier scoped to the SDM Runtime. This is
33 generated by the SDM runtime
34         and is immutable -->
35         <!-- the parent of this instance -->
36         <!-- the host of this instance -->
37         <!-- Name of the member on the parent type -->
38         <!-- Fully qualified type that this is an instance of -->
39         <!-- alias for the id that can be filled in when the instance is
40 created_par
41 member. -->
42         this name must be unique for all instances of the same
43     </xs:complexType>

```

```

1      <xs:complexType name="componentInstance">
2          <xs:complexContent>
3              <xs:extension base="instanceBase">
4                  <xs:sequence>
5                      <xs:element name="portInstances">
6                          <xs:complexType>
7                              <xs:sequence>
8                                  <xs:element
9                                      name="portInstance" type="instanceRef" />
10                                     <!-- the port
11                                     Instances that I own -->
12                                 </xs:sequence>
13                             </xs:complexType>
14                         </xs:element>
15                     </xs:sequence>
16                 </xs:extension>
17             </xs:complexContent>
18         </xs:complexType>
19         <xs:complexType name="compoundComponentInstance">
20             <xs:complexContent>
21                 <xs:extension base="instanceBase">
22                     <xs:sequence>
23                         <xs:element name="portInstances">
24                             <xs:complexType>
25                                 <xs:sequence>
26                                     <xs:element
27                                         name="portInstance" type="instanceRef" />
28                                         <!-- the port
29                                         Instances that I delegate -->
30                                     </xs:sequence>
31                                 </xs:complexType>
32                             </xs:element>
33                         </xs:sequence>
34                     </xs:extension>
35                 </xs:complexContent>
36             </xs:complexType>
37             <xs:element
38                 name="componentInstances">
39                     <xs:complexType>
40                         <xs:sequence>
41                             <xs:element
42                                 name="componentInstance" type="instanceRef" />
43                                 </xs:sequence>
44                             </xs:complexType>
45                         </xs:element>
46                     <xs:element name="wireInstances">
47                         <xs:complexType>

```

```

1      <xs:sequence>
2      <xs:element
name="wireInstance" type="instanceRef" />
3      </xs:sequence>
4      </xs:complexType>
5      </xs:element>
6      </xs:sequence>
7      </xs:extension>
8      </xs:complexContent>
9      </xs:complexType>
10     <xs:complexType name="portInstance">
11     <xs:complexContent>
12     <xs:extension base="instanceBase">
13     <xs:sequence />
14     </xs:extension>
15     </xs:complexContent>
16     </xs:complexType>
17     <xs:complexType name="wireInstance">
18     <xs:complexContent>
19     <xs:extension base="instanceBase">
20     <xs:sequence>
21     <xs:element name="portInstances">
22     <xs:complexType>
23     <xs:sequence>
24     <xs:element
name="portInstance" type="instanceRef" />
25     <!-- the ports that I
have attached -->
26     </xs:sequence>
27     </xs:complexType>
28     </xs:element>
29     </xs:sequence>
30     </xs:extension>
31     </xs:complexContent>
32     </xs:complexType>
33     <xs:complexType name="instanceRef">
34     <xs:attribute name="uniqueId" type="xs:string" />
35     </xs:complexType>
36     <!-- ===== Simple Types
===== -->
37     <xs:simpleType name="fourPartVersionType">
38     <xs:annotation>

```

```

1         <xs:documentation>Four part version numbers where
the segments are in the range 0-65535 </xs:documentation>
2         </xs:annotation>
3         <xs:restriction base="xs:string">
4             <xs:pattern value="(0|[1-5][0-9]{0,4}|[7-9][0-
5             9]{0,3}|6[0-4][0-9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-
6             9]?|655|655[0-2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-
7             9]{0,3}|6[0-4][0-9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-
8             9]?|655|655[0-2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-
9             9]{0,3}|6[0-4][0-9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-
0             9]?|655|655[0-2][0-9]?|655[4-9]|6553[0-5]?)." />
1         </xs:restriction>
2         </xs:simpleType>
3         <xs:simpleType name="publicKeyTokenType">
4             <xs:annotation>
5                 <xs:documentation>Public Key Token: 16 hex digits in
6 size</xs:documentation>
7             </xs:annotation>
8             <xs:restriction base="xs:string">
9                 <xs:pattern value="([0-9]|[a-f]|[A-F]){16}" />
10            </xs:restriction>
11        </xs:simpleType>
12        <xs:attributeGroup name="identity">
13            <xs:attribute name="name" type="xs:string" use="required"
14 />
15            <xs:attribute name="version" type="fourPartVersionType"
16 use="required" />
17            <xs:attribute name="publicKeyToken"
18 type="publicKeyTokenType" use="optional" />
19        </xs:attributeGroup>
20        <xs:simpleType name="guid">
21            <xs:restriction base="xs:string">
22                <xs:pattern value="[0-9a-fA-F]{8}-[0-9a-fA-F]{4}-[0-
23                9a-fA-F]{4}-[0-9a-fA-F]{4}-[0-9a-fA-F]{12}" />
24            </xs:restriction>
25        </xs:simpleType>
</xs:schema>

```

1 The following is an example structure for mappings.

```
2     <?xml version="1.0" encoding="utf-8" ?>
3     <xs:schema targetNamespace="urn:schemas-microsoft-com:sdmMapping"
4     xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:names="urn:schemas-
5     microsoft-com:sdmNames" xmlns="urn:schemas-microsoft-com:sdmMapping"
6     elementFormDefault="qualified" version="0.7" id="sdmMapping">
7         <!-- REVIEW [BassamT]: Do we allow mappings to components
8         within the same compound component? -->
9         <xs:import namespace="urn:schemas-microsoft-com:sdmNames"
10        schemaLocation="SDM7Names.xsd" />
11        <xs:element name="logicalPlacement">
12            <xs:annotation>
13                <xs:documentation>
14                    This file contains the mapping information
15                    between SDM members.
16                    Mappings are constructed in a outside in
17                    fashion, first binding the outer compound component, then its members and so on.
18                </xs:documentation>
19            </xs:annotation>
20            <xs:complexType>
21                <xs:sequence>
22                    <xs:element name="import"
23                    type="names:import" minOccurs="0" maxOccurs="unbounded" />
24                    <xs:element name="placement" minOccurs="0"
25                    maxOccurs="unbounded">
26                        <xs:complexType>
27                            <xs:sequence>
28                                <xs:element
29                                name="memberBinding" type="memberBinding" maxOccurs="unbounded" />
30                                <xs:element
31                                name="wireBinding" type="wireBinding" minOccurs="0"
32                                maxOccurs="unbounded" />
33                            </xs:sequence>
34                            <xs:attribute
35                            name="sourceComponentType" type="xs:string" />
36                            <xs:attribute
37                            name="targetComponentType" type="xs:string" />
38                            <xs:attribute name="name"
39                            type="xs:string" />
40                        </xs:complexType>
41                    </xs:element>
42                </xs:sequence>
43            </xs:complexType>
```



```

1         </xs:element>
2         <!-- a member binding may be a:
3             1. compound component member - in which case we bind all
4             the members and wires of the compound component
5             2. a simple component member - in which case we bind the
6             component and its ports
7             3. a port member - in which case we bind it to a port and there
8             is no further binding
9             -->
10        <xs:complexType name="memberBinding">
11            <xs:sequence>
12                <xs:element name="memberBinding"
13                type="memberBinding" minOccurs="0" maxOccurs="unbounded" />
14                <xs:element name="wireBinding" type="wireBinding"
15                minOccurs="0" maxOccurs="unbounded" />
16            </xs:sequence>
17            <xs:attribute name="sourceMember" type="xs:string"
18            use="required" />
19            <!-- if a target member is not provided then the component
20            must be a compound component and its members
21            will be bound to the members of the compound
22            component that its parent is bound to
23            If a target member is provided and we are binding a
24            compound component, then the ports on the
25            source compound component must be able to be bound to
26            the ports on the target compound component-->
27            <xs:attribute name="targetMember" type="xs:string"
28            use="optional" />
29        </xs:complexType>
30        <!-- wires are bound to a path in the target compound component.
31        This path consists of port, wire and component instances-->
32        <xs:complexType name="wireBinding">
33            <xs:sequence>
34                <xs:element name="path">
35                    <xs:complexType>
36                        <xs:sequence>
37                            <xs:element name="element"
38                            maxOccurs="unbounded">
39                                <xs:complexType>
40                                    <xs:attribute
41                                    name="name" type="xs:string" />
42                                </xs:complexType>
43                            </xs:element>

```

```
1          </xs:sequence>
2          </xs:complexType>
3          </xs:element>
4          </xs:sequence>
5          <xs:attribute name="sourceWire" type="xs:string" />
6          </xs:complexType>
7          <!-- import -->
8          </xs:schema>
```

1 The following is an example structure for names.

```
2     <?xml version="1.0" encoding="UTF-8" ?>
3     <xs:schema targetNamespace="urn:schemas-microsoft-com:sdmNames"
4     xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns="urn:schemas-
5     microsoft-com:sdmNames" elementFormDefault="qualified" version="0.7"
6     id="sdmNames">
7         <xs:import namespace="http://www.w3.org/2001/XMLSchema" />
8         <!-- import creates an alias to another SDM file -->
9         <xs:complexType name="import">
10             <xs:attribute name="alias" type="xs:NCName"
11             use="required" />
12             <xs:attribute name="location" type="xs:NCName"
13             use="optional" />
14             <xs:attributeGroup ref="Identity" />
15         </xs:complexType>
16         <!-- class and type files are identified by name, version and public
17         key -->
18         <xs:attributeGroup name="Identity">
19             <xs:attribute name="name" type="xs:string" use="required"
20             />
21             <xs:attribute name="version" type="fourPartVersionType"
22             use="required" />
23             <xs:attribute name="publicKeyToken"
24             type="publicKeyTokenType" use="optional" />
25         </xs:attributeGroup>
26         <xs:attributeGroup name="namespaceIdentity">
27             <xs:attributeGroup ref="Identity" />
28             <xs:attribute name="signature" type="xs:string"
29             use="optional" />
30             <xs:attribute name="publicKey" type="xs:string"
31             use="optional" />
32         </xs:attributeGroup>
33         <!-- simple version number -->
34         <xs:simpleType name="fourPartVersionType">
35             <xs:annotation>
36                 <xs:documentation>Four part version numbers where
37                 the segments are in the range 0-65535 </xs:documentation>
38             </xs:annotation>
39             <xs:restriction base="xs:string">
40                 <xs:pattern value="(0|[1-5][0-9]){0,4}[[7-9][0-
41                 9]{0,3}|6[0-4][0-9]{0,3}|6[6-9][0-9]{0,2}|65[65[0-4][0-9]{0,2}|65[6-9][0-
42                 9]?|655[655[0-2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}[[7-9][0-
43                 9]{0,3}|6[0-4][0-9]{0,3}|6[6-9][0-9]{0,2}|65[65[0-4][0-9]{0,2}|65[6-9][0-
```

```

1  9]?|655|655[0-2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-
2  9]{0,3}|6[0-4][0-9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-
3  9]?|655|655[0-2][0-9]?|655[4-9]|6553[0-5]?).(0|[1-5][0-9]{0,4}|[7-9][0-
4  9]{0,3}|6[0-4][0-9]{0,3}|6[6-9][0-9]{0,2}|65|65[0-4][0-9]{0,2}|65[6-9][0-
5  9]?|655|655[0-2][0-9]?|655[4-9]|6553[0-5]?)" />
6      </xs:restriction>
7      </xs:simpleType>
8      <!-- public key for verifying signed docs -->
9      <xs:simpleType name="publicKeyTokenType">
10         <xs:annotation>
11             <xs:documentation>Public Key Token: 16 hex digits in
12 size</xs:documentation>
13         </xs:annotation>
14         <xs:restriction base="xs:string">
15             <xs:pattern value="([0-9]|[a-f]|[A-F]){16}" />
16         </xs:restriction>
17     </xs:simpleType>
18 </xs:schema>

```

1 The following is an example structure for settings.

```
2     <?xml version="1.0" encoding="utf-8" ?>
3     <xs:schema targetNamespace="urn:schemas-microsoft-com:sdmSettings"
4     xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns="urn:schemas-
5     microsoft-com:sdmSettings" elementFormDefault="qualified" version="0.7"
6     id="sdmSettings">
7         <xs:import namespace="http://www.w3.org/2001/XMLSchema" />
8         <!-- settings schema, values and constraints -->
9         <xs:complexType name="openSchema">
10             <xs:sequence>
11                 <xs:any namespace="##other" processContents="lax"
12                 />
13             </xs:sequence>
14             </xs:complexType>
15             <xs:complexType name="settingSchema">
16                 <xs:sequence>
17                     <xs:any
18                     namespace="http://www.w3.org/2001/XMLSchema" processContents="skip"
19                     minOccurs="0" maxOccurs="unbounded" />
20                 </xs:sequence>
21                 </xs:complexType>
22                 <xs:complexType name="settingValues">
23                     <xs:sequence>
24                         <xs:any namespace="##other" processContents="lax"
25                         />
26                     </xs:sequence>
27                 </xs:complexType>
28                 <!-- constraints -->
29                 <xs:attributeGroup name="testAttributes">
30                     <xs:attribute name="path" type="xs:string" />
31                     <xs:attribute name="ifNullPath" type="ifNullPath" />
32                     <xs:attribute name="error" type="xs:int" />
33                     <xs:attribute name="errorDesc" type="xs:string" />
34                 </xs:attributeGroup>
35                 <xs:complexType name="simpleTest">
36                     <xs:attributeGroup ref="testAttributes" />
37                 </xs:complexType>
38                 <xs:complexType name="settingConstraints">
39                     <xs:sequence>
40                         <xs:element name="mustExist" type="simpleTest"
41                         minOccurs="0" maxOccurs="unbounded" />
42                         <xs:element name="mustNotExist" type="simpleTest"
43                         minOccurs="0" maxOccurs="unbounded" />
44                     </xs:sequence>
45                 </xs:complexType>
46             </xs:sequence>
47         </xs:complexType>
48     </xs:schema>
```

```

1      <xs:element name="ifExists" type="nestedTest"
2      minOccurs="0" maxOccurs="unbounded" />
3      <xs:element name="ifNotExists" type="nestedTest"
4      minOccurs="0" maxOccurs="unbounded" />
5      </xs:sequence>
6      </xs:complexType>
7      <xs:complexType name="nestedTest">
8      <xs:sequence>
9      <xs:element name="mustExist" type="simpleTest"
10     minOccurs="0" maxOccurs="unbounded" />
11     <xs:element name="mustNotExist" type="simpleTest"
12     minOccurs="0" maxOccurs="unbounded" />
13     <xs:element name="ifExists" type="nestedTest"
14     minOccurs="0" maxOccurs="unbounded" />
15     <xs:element name="ifNotExists" type="nestedTest"
16     minOccurs="0" maxOccurs="unbounded" />
17     </xs:sequence>
18     <xs:attributeGroup ref="testAttributes" />
19     </xs:complexType>
20     <xs:complexType name="deploymentSchema">
21     <xs:sequence>
22     <xs:any namespace="##other" processContents="lax"
23     />
24     </xs:sequence>
25     </xs:complexType>
26     <xs:complexType name="deploymentValues">
27     <xs:sequence>
28     <xs:any namespace="##other" processContents="lax"
29     />
30     </xs:sequence>
31     </xs:complexType>
32     <!-- ===== Simple Types
33     ===== -->
34     <xs:simpleType name="ifNullPath">
35     <xs:restriction base="xs:string">
36     <xs:enumeration value="skip" />
37     <xs:enumeration value="override" />
38     <xs:enumeration value="returnError" />
39     </xs:restriction>
40     </xs:simpleType>
41 </xs:schema>

```

1 The following is an example structure for types.

```
2     <?xml version="1.0" encoding="utf-8" ?>
3     <xs:schema targetNamespace="urn:schemas-microsoft-com:sdmTypes"
4 xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns="urn:schemas-
5 microsoft-com:sdmTypes" xmlns:names="urn:schemas-microsoft-
6 com:sdmNames" xmlns:settings="urn:schemas-microsoft-com:sdmSettings"
7 elementFormDefault="qualified" version="0.7" id="sdmTypes">
8     <xs:import namespace="http://www.w3.org/2001/XMLSchema" />
9     <xs:import namespace="urn:schemas-microsoft-com:sdmSettings"
10 schemaLocation="SDM7Settings.xsd" />
11     <xs:import namespace="urn:schemas-microsoft-com:sdmNames"
12 schemaLocation="SDM7Names.xsd" />
13     <!-- TODO [BassamT]: Normalize the port class refs, port type refs
14 and port members on wire classes, wire types and wire members -->
15     <!-- TODO [BassamT]: Add keys and keyefs for validation -->
16     <!-- TODO [BassamT]: Add support for inlined types -->
17     <!-- TODO [BassamT]: scrub minOccurs and maxOccurs -->
18     <!-- TODO [BassamT]: New name for "class", possibly
19 "deployment" -->
20     <!-- TODO [BassamT]: New name for "host", possibly "provider" --
21 >
22     <!-- REVIEW [BassamT]: Can we merge the definitions of port,
23 component, wire classs in this XSD. It would make it less verbose at the cost more
24 semantic analysis. -->
25     <!-- CONSIDER [BassamT]: General attribute mechanism for things
26 like Singleton, Colocation, Inline. -->
27     <!-- TODO [BassamT]: Bindings: member to component member --
28 >
29     <!-- TODO [geoffo]: ports - are they singleton? -->
30     <!-- TODO [geoffo]: delegation - how do we combine ports? -->
31     <!-- TODO [geoffo] Add back <any> in appropriate places -->
32     <!--
33 =====
34 =====
35 = -->
36     <!-- SDM root element -->
37     <!--
38 =====
39 =====
40 = -->
41     <xs:element name="sdmTypes">
42         <xs:annotation>
```

```

1         <xs:documentation>SDM root element. It is a
2         container for SDM types.</xs:documentation>
3         </xs:annotation>
4         <xs:complexType>
5         <xs:sequence>
6             <xs:element name="import"
7             type="names:import" minOccurs="0" maxOccurs="unbounded" />
8             <xs:element name="information"
9             type="information" minOccurs="0" />
10            <xs:element name="portTypes"
11            minOccurs="0">
12                <xs:complexType>
13                    <xs:sequence>
14                        <xs:element
15                        name="portType" type="portType" minOccurs="0" maxOccurs="unbounded" />
16                    </xs:sequence>
17                </xs:complexType>
18            </xs:element>
19            <xs:element name="componentTypes"
20            minOccurs="0">
21                <xs:complexType>
22                    <xs:sequence>
23                        <xs:element
24                        name="componentType" type="componentType" minOccurs="0"
25                        maxOccurs="unbounded" />
26                    </xs:sequence>
27                </xs:complexType>
28            </xs:element>
29            <xs:element
30            name="compoundComponentType" type="compoundComponentType"
31            minOccurs="0" maxOccurs="unbounded" />
32        </xs:sequence>
33    </xs:complexType>
34    </xs:element>
35    <xs:attributeGroup ref="names:namespaceIdentity" />
36    </xs:complexType>
37    </xs:element>
38    <!-- SDM type library information -->
39    <xs:complexType name="information">
40        <xs:annotation>
41            <xs:documentation>Human readable information
42            about the SDM type library.</xs:documentation>
43        </xs:annotation>
44    </xs:sequence>

```



```

1      minOccurs="0" />      <xs:element name="friendlyName" type="xs:string"
2      minOccurs="0" />      <xs:element name="companyName" type="xs:string"
3      minOccurs="0" />      <xs:element name="copyright" type="xs:string"
4      minOccurs="0" />      <xs:element name="trademark" type="xs:string"
5      minOccurs="0" />      <xs:element name="description" type="xs:string"
6      minOccurs="0" />      <xs:element name="comments" type="xs:string"
7      minOccurs="0" />
8          </xs:sequence>
9          </xs:complexType>
10         <!--
11         =====_u61
12         ?=====
13         == -->
14         <!-- base complexType for component, port, and wire types -->
15         <!--
16         =====
17         = -->
18         <xs:complexType name="baseType">
19             <xs:annotation>
20                 <xs:documentation>base type for Component Type
21                 and Port Type.</xs:documentation>
22             </xs:annotation>
23             <xs:sequence>
24                 <xs:element name="deployment"
25                 type="settings:deploymentValues" minOccurs="0" />
26                 <xs:element name="settings"
27                 type="settings:settingValues" minOccurs="0" />
28                 <xs:element name="settingSchema"
29                 type="settings:settingSchema" minOccurs="0" />
30                 <xs:element name="hostConstraints"
31                 type="hostConstraints" minOccurs="0" />
32                 <xs:element name="hostedClasses"
33                 type="hostedClassesList" minOccurs="0" />
34                 <!-- deployment section contains the deployment
35                 instructions for the type. The
36                 schema for this deployment is specified on the
37                 "unit". -->

```

```

1      terms of the schema      <!-- The settings for this component. These are in
2                                that is specified on the unit. -->
3      exposed by the type.      <!-- Setting Schema. New setting schema that can be
4                                The values for this schema are set on members of
5      this type. Note           also that we will support flowing setting values to
6                                on the unit. -->
7      settings of the host that <!-- This section contains any constraints on the
8                                is hosting this unit. Note that there can be more than
9                                each unit and we can supply constraints on each. -->
10     any constraints on the classs' settings. <!-- This section contains the list of hosted units and
11     schema. Note that there can be           The constraints are specified in terms of the class's
12     more than one class that is hosted here. -->
13     </xs:sequence>
14     <xs:attribute name="class" type="xs:string" use="required"
15     />
16     <xs:attribute name="name" type="xs:string" use="required"
17     />
18     </xs:complexType>
19     <!--
20     =====
21     =====
22     = -->
23     <!-- Constraints -->
24     <!--
25     =====
26     =====
27     = -->
28     <xs:complexType name="hostConstraints">
29         <xs:annotation>
30             <xs:documentation>Host constraints are constraints
31             against the classes that can host this component or port type.</xs:documentation>
32         </xs:annotation>
33         <xs:sequence>
34             <xs:element name="hostConstraint">
35                 <xs:complexType>

```

```

1          <xs:sequence>
2      type="settings:settingConstraints" />
3          </xs:sequence>
4          <xs:attribute name="host"
5      type="xs:string" use="required" />
6          <!-- the name of the component unit that
7      is the host.
8          May include an alias, for
9      example, "otherSDM:myPortType" -->
10         </xs:complexType>
11         </xs:element>
12         </xs:sequence>
13     </xs:complexType>
14     <xs:complexType name="hostedClassesList">
15         <xs:annotation>
16             <xs:documentation>These are constraint against the
17 classes that this type can host</xs:documentation>
18         </xs:annotation>
19         <xs:sequence>
20             <xs:element name="hostedClass"
21 maxOccurs="unbounded">
22                 <xs:complexType>
23                     <xs:sequence>
24                         <xs:element name="constraints"
25 type="settings:settingConstraints" />
26                     </xs:sequence>
27                     <xs:attribute name="class"
28 type="xs:string" use="required" />
29                     <!-- the name of the component class
30 that we are hosting.
31                     May include an alias, for
32 example, "otherSDM:myPortType" -->
33                         </xs:complexType>
34                     </xs:element>
35                 </xs:sequence>
36             </xs:complexType>
37         </xs:sequence>
38     </xs:complexType>
39     <!--
40 =====
41 =====
42 =====
43 = -->
44     <!-- Ports -->
45

```

```

1      <!--
2      =====
3      =====
4      = -->
5      <xs:complexType name="portType">
6          <xs:complexContent>
7              <xs:extension base="baseType">
8                  <xs:sequence>
9                      <xs:element name="portConstraints"
10 type="portConstraints" minOccurs="0" maxOccurs="unbounded" />
11                  </xs:sequence>
12              </xs:extension>
13          </xs:complexContent>
14      </xs:complexType>
15      <xs:complexType name="portConstraints">
16          <xs:annotation>
17              <xs:documentation>These are constraint against the
18 classes that this type can host</xs:documentation>
19          </xs:annotation>
20          <xs:sequence>
21              <xs:element name="portConstraint">
22                  <xs:complexType>
23                      <xs:sequence>
24                          <xs:element name="constraints"
25 type="settings:settingConstraints" />
26                      </xs:sequence>
27                      <xs:attribute name="portClass"
28 type="xs:string" use="required" />
29                      <xs:attribute name="portType"
30 type="xs:string" use="optional" />
31                      <xs:attribute name="minOccurs"
32 type="xs:int" use="optional" />
33                      <xs:attribute name="maxOccurs"
34 type="xs:int" use="optional" />
35                      <xs:attribute name="visible"
36 type="xs:boolean" use="optional" />
37                  <!-- port type is here to allow you to bind
38 to a port type that has behavior not exposed
39                      through settings. In this case settings
40 would not be enough to allow an appropriate port to be identified -->
41                  <!-- visible identifies whether the
42 application wants to see ports that match this constraint. An application may

```

```

1         only choose to constrain port that it is
connected to without wanting connection values for their endpoints -->
2         </xs:complexType>
3         </xs:element>
4         <!-- we probable need a mechanism here to allow a
port to identify sets of optional port endpoints.
i.e. I can bind to X,Y or Z but I must bind to at least
one-->
5         </xs:sequence>
6         </xs:complexType>
7         <!--
=====
=====
8     = -->
9         <!-- Components -->
10        <!--
=====
=====
11    = -->
12        <xs:complexType name="componentType">
13            <xs:complexContent>
14                <xs:extension base="baseType">
15                    <xs:sequence>
16                        <xs:element name="ports"
17                        type="portsList" minOccurs="0" />
18                    </xs:sequence>
19                    <!-- the exclusive attribute indicates whether an
instance of the type at the layer below can host instance of other types
- shallow implies that the only instances hosted by the
instance are instance of this type
- deep implies that the host must also be marked as
exclusive -->
20                    <xs:attribute name="exclusive" type="depth"
use="optional" default="notSet" />
21                </xs:extension>
22            </xs:complexContent>
23        </xs:complexType>
24        <xs:complexType name="portsList">
25            <xs:sequence>
                <xs:element name="port" minOccurs="0"
maxOccurs="unbounded">
                    <xs:complexType>
                        <xs:sequence>

```

```

1         type="settings:settingValues" minOccurs="0" />
2         <!-- member setting values. These
3         are in terms of the setting schema
4         specified on the type -->
5         </xs:sequence>
6         <xs:attribute name="name"
7         type="xs:string" use="required" />
8         <xs:attribute name="type"
9         type="xs:string" />
10        </xs:complexType>
11        </xs:element>
12        </xs:sequence>
13        </xs:complexType>
14        <!--
15        =====
16        =====
17        = -->
18        <!-- Compound Component Types. -->
19        <!--
20        =====
21        =====
22        = -->
23        <xs:complexType name="compoundComponentType">
24            <xs:sequence>
25                <xs:element name="components" type="components"
26                minOccurs="0" />
27                <xs:element name="wires" type="wires"
28                minOccurs="0" />
29                <xs:element name="delegatePorts"
30                type="delegatePorts" minOccurs="0" />
31                <xs:element name="delegateHostedClasses"
32                type="delegateHostedClasses" minOccurs="0" />
33                <!-- delegate ports -->
34                <!-- delegate hosts. These allow a compound
35                component to act as a host just as a simple
36                component. -->
37            </xs:sequence>
38            <xs:attribute name="name" type="xs:string" use="required"
39            />
40            <!-- shallow co-location implies that instances of the
41            members of this compound component will be placed on the

```

```

1         same host instance at the layer below - deep co-
location implies that that host is also colocated-->
2         <xs:attribute name="colocate" type="depth" use="optional"
default="notSet" />
3         <!-- the open attribute indicates whether bindings can see the
internal structure of the compound component
4         or simply bind to the compound component as though it
was a simple component -->
5         <xs:attribute name="open" type="xs:boolean" use="optional"
default="false" />
6         <!-- the exclusive attribute indicates whether an instance of
the type at the layer below can host instance of other types
7         - shallow implies that the only instances hosted by the
instance are instance of this type
8         - deep implies that the host must also be marked as
9         exclusive -->
10        <xs:attribute name="exclusive" type="depth" use="optional"
default="notSet" />
11        </xs:complexType>
12        <xs:complexType name="components">
13        <xs:sequence>
14        <xs:element name="component" minOccurs="0"
maxOccurs="unbounded">
15        <xs:complexType>
16        <xs:sequence>
17        <xs:element name="settings"
type="settings:settingValues" minOccurs="0" />
18        <!-- member setting values. These
are in terms of the setting schema
19        specified on the type -->
20        </xs:sequence>
21        <xs:attribute name="name"
type="xs:string" use="required" />
22        <xs:attribute name="type"
type="xs:string" use="required" />
23        <xs:attribute name="singleton"
type="xs:boolean" use="optional" default="false" />
24        </xs:complexType>
25        </xs:element>
</xs:sequence>
</xs:complexType>
<xs:complexType name="wires">
<xs:sequence>

```

```

1      <xs:element name="wire" minOccurs="0"
maxOccurs="unbounded">
2          <xs:complexType>
3              <xs:sequence>
4                  <xs:element name="members">
5                      <xs:complexType>
6                          <xs:sequence>
7                              <xs:element
name="member" type="componentPortRef" minOccurs="0"
maxOccurs="unbounded" />
8                          </xs:sequence>
9                      </xs:complexType>
10                     </xs:element>
11                     <!-- member setting values. These
are in terms of the setting schema
12                     specified on the type -->
13                     </xs:sequence>
14                     <xs:attribute name="name"
type="xs:string" use="required" />
15                     <xs:attribute name="protocol"
type="xs:string" />
16                 </xs:complexType>
17             </xs:element>
18         </xs:sequence>
19     </xs:complexType>
20     <xs:complexType name="delegatePorts">
21         <xs:sequence>
22             <xs:element name="delegatePort" minOccurs="0"
maxOccurs="unbounded">
23                 <xs:complexType>
24                     <xs:attribute name="name"
type="xs:string" />
25                     <xs:attribute name="componentName"
type="xs:string" />
26                     <xs:attribute name="portName"
type="xs:string" use="optional" />
27                 </xs:complexType>
28             </xs:element>
29         </xs:sequence>
30     </xs:complexType>
31     <xs:complexType name="componentPortRef">
32         <xs:attribute name="componentName" type="xs:string" />

```



```

1      <xs:attribute name="portName" type="xs:string"
    use="required" />
2      </xs:complexType>
3      <xs:complexType name="delegateHostedClasses">
4          <xs:sequence>
5              <xs:element name="hostedClassRef"
6                  maxOccurs="unbounded">
7                      <xs:complexType>
8                          <xs:attribute name="componentName"
9                              type="xs:string" />
10                         <xs:attribute name="hostedClass"
11                             type="xs:string" use="required" />
12                     </xs:complexType>
13                 </xs:element>
14             </xs:sequence>
15         </xs:complexType>
16     <!--
17     =====
18     =====
19     = -->
20
21     <!-- SimpleTypes. -->
22     <!--
23
24     =====
25     =====
26     = -->
27
28     <!-- depth identifies whether the settings applies to only the next
29     layer (shallow), or must apply to the next layer and be set on the next layer (deep)
30     -->
31
32     <xs:simpleType name="depth">
33         <xs:restriction base="xs:string">
34             <xs:enumeration value="notSet" />
35             <xs:enumeration value="shallow" />
36             <xs:enumeration value="deep" />
37         </xs:restriction>
38     </xs:simpleType>
39 </xs:schema>

```

SDM Runtime

The SDM Runtime (or just runtime) hosts an implementation of the SDM. It is a highly available distributed service that exposes a set APIs for manipulating the SDM type, member and instance space. The runtime is responsible for tracking all SDM instances in a consistent manner. It provides machinery for deployment, versioning, security and recovery. Fig. 27 represents the logical architecture of the SDM runtime.

The SDM runtime consists of the following:

- SDM Runtime – this is the SDM Runtime implementation. It is a distributed implementation that will run on one or more physical machines. The runtime exposes its functionality through the SDM API which is set of calls that manipulate the SDM and instances.
- SDM Store – this is a durable store for SDM Models and instances. This store is highly available and its consistency is critical. This store will survive catastrophic events.
- Service Deployment Units – this is a read-only store for SDUs. Just like the SDM store it is highly available and will survive catastrophic events.
- Component Implementation Host – this is framework for hosting the CLR code that is referenced from SDM components.

The SDM Runtime is typically used by the following client classes:

- Component Instances – these are component instances that communicate with the runtime using the SDM Runtime Library (RTL). We distinguish between

1 two types of component instances – runtime-hosted component instances and
2 non runtime-hosted component instances.

- 3 • Development and Deployment tools – these include the SDM compiler, SDU
4 installation tools as well as other development tools.
- 5 • Management tools – these are privileged tools that are used for administering
6 and managing the runtime itself.

7
8 Clients communicate with the runtime through the SDM Runtime Library (RTL).

9 They typically perform operations that include:

- 10 • Installing / Uninstalling SDUs: This is the process of adding and removing new
11 SDUs into a running instance of the SDM Runtime.
- 12 • Adding, removing and modifying SDM types and instances: clients can create
13 new components, ports and wire types.
- 14 • Creating and deleting instances: clients can create new components, port and
15 wire instances.
- 16 • Sourcing and sinking events: when changes are made to the type and/or
17 instance space the runtime will send events the affected clients. Events can also
18 be triggered on specific operations such as setting the port binding information.
- 19 • Query the type and instance space: Clients can reflect on the type and instance
20 space.

21 22 Service Definition Model Runtime Architecture

23 Introduction

24 This document discusses the Service Definition Model (SDM) and SDM
25 Runtime. A technical discussion of the runtime architecture, core features and

1 implementation are provided. The intended audience is technical evaluators of
2 BIG, developers who intend to author services and components, or others with an
3 interest in the details of the system.

4 Services Era

5 Over the last decade we have witnessed the Internet emerge as a computing
6 platform. More and more software companies are adopting the “software as a
7 service” model. Services are typically comprised of several components running
8 on many machines including servers, networking gear and other specialized
9 hardware. Loosely coupled, asynchronous programming models are becoming the
10 norm. Scalability, availability and reliability are critical to the success of such
11 services.

12 We are also witnessing a change in hardware trends. High density servers
13 and specialized network hardware are widespread in data centers. Switched fabrics
14 are replacing system buses providing greater flexibility in system configurations.
15 Hardware cost plays a small role in the Total Cost of Ownership metric. This has
16 been replaced by the cost of maintaining a dedicated operations staff. Rock-solid
17 operational practices are rare but absolutely vital to any service. These practices,
18 for the most part, are implemented by people.

19 Effectively the focus of development is shifting from the single PC into the
20 network of PCs. Yet with all these changes have come a plethora of new problems
21 for service developers, software vendors, hardware vendors, and end-users:

- 22 ■ Services are large and complex – they are time-consuming to
23 develop, difficult and costly to maintain, and risky to extend with
24 additional functionality.
25

- 1 ■ Services are monolithic – they rely on custom components and
2 configurations. Portions of the service cannot be removed, upgraded
3 independently, or replaced with alternatives.
- 4 ■ Services rely on specific hardware configurations – whether it's a
5 certain network topology or a dependency on a specific network
6 appliance machine. This significantly reduces the ability to host a
7 service in a different environment.
- 8 ■ Services are developed in silos – due to the lack of a common
9 platform, sharing code or even best operational practices is a
10 daunting task.
- 11 ■ Operational nightmare – most services require a staff of operations
12 personnel to function. The operations staff must be trained in the
13 specifics of each service and retrained as the service evolves.
14

15
16 Some of these problems are not unlike those during the DOS era (circa
17 1980's). DOS defined valuable core services for application developers such as
18 disk management, file system, console facilities, etc. It did, however, leave many
19 complex tasks up to the ISVs. As an example, WordPerfect and Lotus 123 both
20 independently had to write printer drivers in order to support printing within their
21 respective applications. Similarly printer hardware vendors had to make deals with
22 the software companies in order to have a successful product. The barrier to entry
23 for writing a DOS application and hardware vendors was exceptionally large. This
24 resulted in only a few successful software companies.
25

1 Windows addressed this problem by defining a platform which dramatically
2 reduced the barrier to entry. Windows defined an abstraction layer for most
3 hardware on the PC platform. This relieved the developers from having to worry
4 about supporting specific hardware devices. Windows managed all resources
5 within the PC including memory, disk and network. It also came with a wealth of
6 services that can be utilized by application developers. This platform sparked
7 enormous growth in the industry. Software vendors that targeted the Windows
8 platform were extremely productive. Many new hardware vendors emerged with
9 cheaper hardware due to the commoditization effect of Windows.

10 The services era has yet to experience such growth -- the revolution that has
11 happened on the desktop machine needs to happen with services.

12 13 BIG services Platform

14
15 BIG is creating a platform for highly available and scalable services. This
16 platform will enable:

- 17 ■ Development of distributed, scalable and highly available services
18 using Visual Studio and reusable building blocks like SQL, IIS, etc.
- 19 ■ Deployment across a set of abstracted hardware and software
20 resources which are automatically allocated, purposed and
21 configured.
- 22
23 ■ Lowering the cost of ownership through automation of operational
24 best practices.
- 25

- Procurement of standardized data center hardware that leverages commodity economics.

The BIG platform is an extension to the Windows platform and builds on the existing technologies such as .NET, SQL Server and other Microsoft assets.

The BIG services platform is comprised of many pieces, including:

- Hardware reference platform that aggregates commodity hardware to build a single large computer that we call the BIG Computer. This includes many interconnected servers, network devices, and storage.
- Hardware abstraction layer that virtualizes resources. Enables dynamic hardware binding and re-deployment and automated network configuration
- Service Definition Model (SDM) for developers to describe an entire service. Enables developers to rapidly build new services using highly available SQL, IIS and other reusable building block components
- Highly available runtime that supports the SDM. Enables hosting multiple scalable services inside the BIG Computer.
- Operations logic framework for automating operational best practices. Enables policy expression and enforcement

This document will focus solely on the SDM and the SDM Runtime.

Service Definition Mode

1 This section will discuss the Service Definition Model (SDM). Please refer
2 to the “Service Definition Model Language” document for a complete technical
3 description of the SDM and the SDML language.

4 The SDM is the foundation on which all services are built. The SDM:

- 5 ■ Enables the composition of services from smaller units. These units
6 form the basis of hardware and software abstraction.
- 7 ■ Serves as a live blueprint of a service – the SDM captures the overall
8 structure of a service in a scale-invariant manner.
- 9 ■ Provides a framework for automating operational practices and
10 promotes their reuse.
- 11 ■ Defines standards for deployment, reuse, discovery, versioning, and
12 recovery of services.
- 13
- 14

15 Component Model for Services

16 In essence, the SDM is a component model for services. Like traditional
17 component models, the SDM defines primitives on which more complex
18 functionality can be built. Let’s consider an analogy; Microsoft’s Component
19 Object Model (COM) defined a programming model for authoring components. It
20 standardized on how components are packaged, registered, activated, discovered
21 etc. COM mandated strict rules related to lifetime, memory management, and
22 interface implementation. These primitives were essential for interoperability – it
23 allowed components to be treated as black boxes. Com was the basis for more
24 sophisticated services such as persistent storage, eventing, automation and OLE.
25

1 The SDM is defining a component model for services. This model is well
2 suited for loosely coupled, distributed and asynchronous services. The SDM
3 defines standards for deployment, versioning, recovery and scoping. The SDM is
4 the model in which more sophisticated services such as network management,
5 hardware management, storage abstraction, etc. are delivered.

6 How does the SDM compare to other component models?

7 Certainly technologies such as DCOM and CORBA among others have well
8 defined methods for developing applications based on reusable components.

9 However, while existing component technologies are powerful, they have not been
10 widely successful in the Internet or loosely coupled scenarios. This is largely due
11 to the following:

- 12 ■ Existing component technologies are not designed for the large scale
13 – most implementations are optimized for a single machine or a
14 small number of machines. Internet applications typically involve
15 many interrelated components running on many machines.
- 16
17 ■ Existing component technologies mandate invocation protocols such
18 as RPC – they do not leverage well-established network protocols
19 nor do they allow diverging protocols.
- 20
21 ■ Existing component technologies lack a concept of an application –
22 most have well developed definitions of components but lack an
23 overall definition of an application that is composed of smaller
24 components.
25

- Existing component technologies are limited to software running on a general purpose computer – single-purpose network devices can not participate as components.

That said there is a lot of thinking that has gone into existing component technologies that is still significantly relevant to the services world.

SDM Fundamentals

The SDM is a declarative definition of the structure of a service. This definition is in terms of components, ports, and wires:

- Components are units of implementation, deployment and operations. Components can be a dedicated server running .NET Server, an IIS virtual web site on a shared machine or network appliance such as a Cisco LocalDirector. Components expose functionality through ports and establish communicating paths through wires. Components can be nested within outer components are referred to as compound components.
- Ports are named endpoints that have an associated type. Port types often represent a protocol, for example, HTTP server. Ports capture the information required for establishing communication.
- Wires are the permissible communication paths between ports. They declare the topological relationship between ports.

1 Services are authored using a declarative Service Definition Model Language

2 (SDML). Let's consider an example:

```
3     using System;
4     using System.Iis;
5     using System.Sql;
6
7     [sdmassembly:name("MyService")];
8     [sdmassembly:version(1)];
9
10    componenttype MyFrontEnd : AspApplication
11    {
12        port SqlClient catalog;
13        implementation "MyFE, MyClrAssembly";
14    }
15
16    componenttype MyBackEnd : SqlDatabase
17    {
18        implementation "MyBE, MyClrAssembly";
19    }
20
21    componenttype MyService
22    {
23        component MyFrontEnd fe;
24        component MyBackEnd be;
25
26        port http = fe.http;
27
28        wire SqlTds tds
29        {
30            fe.catalog;
31            be.sqlServer;
32        }
33
34        implementation "MyService, MyClrAssembly";
35    }
```

22 As can be seen the syntax for SDML borrows heavily from C#. SDML defines
23 component, port and wire types. If we walk through this definition:

- 24 ■ The using directive references namespaces of SDM types. These
25 include the system namespace which is provided by the SDM

1 runtime and defines basic types such as the http wire type. The other
2 namespaces define types associated with IIS and SQL Server.

- 3 ■ The assembly name and assembly version provide a strong name for
4 the SDM assembly. Note that this is nothing to do with a CLR
5 assembly. An SDM assembly is the smallest unit of SDM
6 deployment. It is named and contains a collection of component,
7 port and wire types. SDM assemblies should not be confused with
8 CLR assemblies – they are completely distinct.
- 9
10 ■ A componenttype called MyFrontEnd is declared that inherits from
11 the component type AspApplication which is a referenced type
12 defined in the System.Iis SDM assembly. Components are
13 abstractions; they refer to a class and not instances. MyFrontEnd
14 identifies a component from which zero or more component
15 instances can be created.
- 16
17 ■ port SqlClient catalog; declares a port on the MyFrontEnd
18 component of type SqlClient. The port is called "catalog". This port
19 is in addition to the ports, components and wires that MyFrontEnd
20 inherits from base component type AspApplication.
- 21
22 ■ The implementation keyword references an implementation for the
23 component type. This implementation is a reference to a CLR class
24 within a CLR assembly. This can be thought of as an entry point or
25 constructor for the component type. When a component instance is
created this code is invoked.

- 1 ■ The MyService component type is defined with two sub-components
2 called fe and be. These are of type MyFrontEnd and MyBackEnd.
3 Instances of component MyService can subsequently have instances
4 of fe and be forming a hierarchy of component instances.
- 5 ■ port http = fe.http; declares a port on the MyService component type
6 that is delegated to the http port on the fe component.
- 7 ■ wire SqlTds tds declares a wire in the MyService component type of
8 type SqlTds, with the name tds. Two ports are attached to the wire.
9 This declaration means an instance of MyService can have zero or
10 more instances of wire tds and each of those wire instances can have
11 catalog ports from fe components and sql ports from be components
12 attached to them.
13

14 It is often helpful to consider a graphical representation of services. See Fig. 28.
15 Boxes represent components, diamonds represent ports and lines represent wires.

17 Component Implementation

18 Every component can reference an implementation in the form of a CLR
19 class within a CLR assembly. The CLR assembly is hosted by the SDM Runtime
20 and will be invoked at component instantiation time. The CLR class that
21 implements the SDM component can perform SDM operations by calling the
22 SDM Runtime API. This will be described in great detail later in this document.
23 The following is a C# code snippet for the implementation of the MyService SDM
24 component type from above.

25 using System;

1 using Microsoft.SDM;

2 public class MyService: SDMComponentInstance

3 {

4 public override OnCreate(...)

5 {

6 SDMComponent fe1 = CreateComponentInstance("fe", "");

7 SDMComponent fe2 = CreateComponentInstance("fe", "");

8 SDMComponent be1 = CreateComponentInstance("be", "");

9 SDMWire tds1 = CreateWire instanceance("tds");

10 tds1.Members.Add(fe1.Ports["catalog"]);

11 tds1.Members.Add(fe2.Ports["catalog"]);

12 tds1.Members.Add(be1.Ports["sqlServer"]);

13 }

14 }

15 This code defines a C# class MyService that inherits from the SDMComponent.

16 The class overrides the OnCreate() method and creates two instances of the fe components, one instance of the be component and one wire instance. It then adds three ports to the wire instance.

17 This CLR code is compiled into an assembly called MyClrAssembly that is
18 referenced within the SDM for MyService. When a component of type MyService
19 is instantiated this code will be invoked and the OnCreate method will be called.

20 [BassamT] Consider showing the strongly-typed version of the C#
21 code.
22
23
24
25

Instances

SDML is used to define component, port and wire types; it does not define instances. Instances can be created using the SDM Runtime API as we saw in the C# code above. The C# code above created a number of instances and formed a wiring topology in the instance space. These instances will be tracked by the SDM Runtime. For example the SDM Runtime will store the following information after the OnCreate call completes above:

```
component instance ms[1]
    port instance http[1]
    component instance fe[1]
    component instance fe[2]
    component instance be[1]
    wire instance tds[1]
        fe[1].catalog
        fe[2].catalog
        be[1].SqlServer;
```

NOTE: The syntax used here is not SDML; it is used to illustrate the instance space that is tracked by the SDM runtime.

ms[1] is a component instance that has three children component instances fe[1], fe[2] and be[1]. fe[1] and fe[2] are instance of the fe component. be[1] is an instance of the be component. tds[1] is a wire instance that contains three members. Graphically, the instance space shown in Fig. 29.

Components instances have real physical manifestations – fe[1] and fe[2] in this example are two ASP.NET applications that are running on IIS running on a Windows machine. When the call to CreateComponentInstance was made a new ASP.NET application was created and configured on an IIS box. A number of

1 intermediate steps could have also been invoked – for example, the caller’s credit
2 card has been charged for using the new resource or a new machine has been
3 allocated due to lack of capacity. Later in this document we will examine the
4 machinery behind component instantiation.

5 6 Service Deployment Units

7
8 The SDM model for MyService defined the structure of the service in terms
9 of component, ports and wires. This resulted in an SDM Assembly that can be
10 installed on an SDM Runtime machine. Obviously, the SDM assembly is not
11 enough for instantiating the service. In addition to the SDM assembly we must
12 also consider the CLR assemblies that are the implementations of components. We
13 must also consider the ASP.NET code, SQL scripts and whatever else is needed by
14 the service. The sum of all these pieces is packaged up into a Service Deployment
15 Unit (or SDU). See Fig. 30.

16 17 SDM Runtime

18 The SDM Runtime (or just runtime) hosts an implementation of the SDM.
19 It is a highly available distributed service that exposes a set APIs for manipulating
20 the SDM type, member and instance space. The runtime is responsible for tracking
21 all SDM instances in a consistent manner. It provides machinery for deployment,
22 versioning, security and recovery.

23 This section describes the design and implementation of the SDM Runtime
24 as proposed for the BIG V1.0 release. While there can certainly be different
25 embodiments of the SDM Runtime we will focus on one throughout this document

1 – the highly available SDM Runtime implementation that will be hosted on the
2 BIG Computer (see. _____ for more details).

4 Runtime Architecture

5
6 Fig. 27 represents the logical architecture of the SDM runtime.

7 The SDM runtime consists of the following:

- 8 ■ SDM Runtime – this is the SDM Runtime implementation. It is a
9 distributed implementation that will run on one or more physical
10 machines. The runtime exposes its functionality through the SDM
11 API which is set of calls that manipulate the SDM and instances.
- 12 ■ SDM Store – this is a durable store for SDM Models and instances.
13 This store is highly available and its consistency is critical. This
14 store will survive catastrophic events.
- 15 ■ Service Deployment Units – this is a read-only store for SDUs. Just
16 like the SDM store it is highly available and will survive
17 catastrophic events.
- 18 ■ Component Implementation Host – this is framework for hosting the
19 CLR code that is referenced from SDM components.

20
21 The SDM Runtime is typically used by the following client classes:

- 22 ■ Component Instances – these are component instances that
23 communicate with the runtime using the SDM Runtime Library
24
25

1 (RTL). We distinguish between two types of component instances –
2 runtime-hosted component instances and non runtime-hosted
3 component instances.

- 4 ■ Development and Deployment tools – these include the SDM
5 compiler, SDU installation tools as well as other development tools.
6
- 7 ■ Management tools – these are privileged tools that are used for
8 administering and managing the runtime itself.

9 Clients communicate with the runtime through the SDM Runtime Library (RTL).

10 They typically perform operations that include:

- 11 ■ Installing / Uninstalling SDUs: This is the process of adding and
12 removing new SDUs into a running instance of the SDM Runtime.
13
- 14 ■ Adding, removing and modifying SDM types and instances: clients
15 can create new components, ports and wire types.
- 16 ■ Creating and deleting instances: clients can create new components,
17 port and wire instances.
18
- 19 ■ Sourcing and sinking events: when changes are made to the type
20 and/or instance space the runtime will send events the affected
21 clients. Events can also be triggered on specific operations such as
22 setting the port binding information.
- 23 ■ Query the type and instance space: Clients can reflect on the type
24 and instance space.
25

Type, member and instance space

The relationship between a component type, component and component instance is analogous to class, class member and object in modern object-oriented languages. SDM defines a separation between the type, member and instance space. Component types are in the type space, components are in the member space and component instances are in the instance space. Fig. 31 illustrates the separation between the three spaces.

The “member space” contains instances of the type space. The “instance space” contains instances of the member space. The SDM Runtime is responsible for tracking all three spaces and the relationship between them. This information is stored within the SDM store and can be queried by using the Runtime API. Components and wires can have zero or more instances. Ports can only have one instance.

The SDM member and instance space conform to a strict hierarchy. All components within the member and instance space are arranged in a tree. The root component is a special component referred to as the “root” or “universal” components. Let’s look at the member tree from the MyService example in the previous section (Fig. 32). The boxes represent components and the lines are parent/child relationships. myService is a member component of the root component. The instance tree might look as shown in Fig. 33. Notice that there are two instance of the myService component with a different number of children instances. myService[1].fe[1] and myService[2].fe[1] have same component member “fe” and have the same component type “MyFrontEnd” but otherwise are

1 completely distinct component instances. “root[1]” is the only instance of the root
2 component.

3 4 Component Instantiation

5
6 One of the fundamental operations provided by the SDM runtime is
7 component instantiation. This is the process in which a component instance comes
8 into existence. Unlike traditional component models where creating an instance
9 (or an object) typically involves allocating and initializing a chunk of memory for
10 the instance, SDM components typically involve many steps performed by
11 different parties and can take hours if not days to complete. For example, when a
12 component of type ASP.NET application is instantiated the result is a new virtual
13 web site on a machine running IIS followed by a configuration act. Consider a
14 scenario where the capacity on the IIS machines has been reached and a new one
15 has to be allocated before an ASP.NET application is instantiated. This process
16 might take hours as it will involve allocating a new machine from a pool, possibly
17 incurring a billing charge, and installing the operating system including IIS.
18 The SDM Runtime supports two ways to instantiate components 1) Factory
19 instantiated components and 2) runtime instantiated components. These methods
20 are discussed briefly below. Please refer to the “Component Instantiation”
21 specification for more details.
22
23
24
25

Factory instantiated components

Component factories (or just factories) are the entities responsible for creating instances for one or more component types. Factories are themselves components that expose one or more ports for the purposes of instantiation. One way to think of factories is as resource managers. The resource they are managing is the component type. Factories know how to map a resource into an instance of a component. For example, assume we had a component of type “File Storage”. When this component is instantiated an NTFS directory will be created and appropriate ACLs will be provisioned. The factory for this component might manage a number of Windows machines for the purpose of allocating storage. The factory is responsible for creating the NTFS share, setting the ACLs, quotas etc. Component factories play an important role in the SDM Runtime. Since they typically are managing resources on behalf of services they are expected to be reliable and highly available. While the number of component factories supported by the SDM runtime is open ended we expect the BIG V1.0 will have a small number of base component factories. They are:

- Hardware – this is base level factory that is responsible for allocating instances of hardware and managing them. For example, it can allocate a server machine with 1GB of memory, or a storage device such as NAS.
- Network – this factory responsible for VLANs, public IP addresses, DNS names etc.

- 1 ■ PC – this factory can allocate a machine and deploy a full OS-image
- 2 on it.
- 3 ■ Storage – this factory is responsible for managing and allocating
- 4 storage.
- 5 ■ Software resources – such as ASP.NET, IIS Web Site, SQL Server
- 6 Database etc.

8 Instantiation process

9 Factories must register with the SDM runtime specifying which component
10 types they are responsible for creating instances of. At a high level, the process of
11 instantiation is as follows:

12 The caller asks the SDM runtime for the component factory for a given
13 component type.

- 14 1. The SDM runtime is responsible for finding the appropriate
- 15 component factory and returning it to the caller.
- 16 2. The caller then communicates with the component factory directly
- 17 and asks it to create one or more instances.

18 Running Factory Table

19 The SDM runtime will maintain a table of the component types and their
20 appropriate factories. Every component instance has a running factory table. The
21 running factory table structure is as follows:

22 (ComponentTypeID, PortType) -> (PortInstance, [cookie])

1 Component instances can add/remove entries in their tables as well as any of their
2 direct children's tables. By default, the running factory table of the parent is
3 inherited when a new child component instance is created.

4 The running factory table is tracked for every component instance in order
5 to support different factories for the same component type in different contexts.
6 Since factories are typically where resources are allocated hosting environments
7 might mandate different policies for resource allocation. For example, consider a
8 scenario where a hosting entity such as Digex has different plans for their
9 customers. Customers that paid for Gold will get a dedicated IIS box and
10 customers that paid for Silver will get a shared IIS box. The customer's service
11 contains a component of type "ASP.NET application" and it is unaware of whether
12 it will be hosted on a dedicated IIS machine or a shared one. Digex might
13 implement this as shown in Fig. 34.

14 Digex is a component that has two component factories Gold Factory and
15 Silver Factory. The factories are components themselves. Digex also defines to
16 other components called "Gold" and "Silver". These "Gold" components will be
17 the parent of all services that have paid for the Gold Service.

18 When Digex is instantiated it will create an instance of the factories and
19 also instances of the "Gold" and "Silver" components. Gold[1] will have its own
20 running factory table. Digex will register the Gold factory in this table by calling
21 the appropriate SDM runtime API. When a new customer's service is instantiated
22 as a child of Gold[1] and it will inherit the running factory table of Gold[1]. This
23 means that when a component instance of "ASP.NET application" is created the
24 Gold Factory will handle this request and charge the customer's account
25 appropriately.

Factory tracking

The SDM runtime will keep track of the factory that created each component instance. See Fig. 35. The dotted lines represent a “created by” relationship between a component instance and the factory that created it. As mentioned above the factories are components themselves and therefore they must have factories. To end the infinite recursion the runtime will be the factory for a “runtime-hosted components” as described below. Note also that the root component instance is special and it is its own factory.

Factories and Transactions

Factories will support transactions to relieve service developers from having to worry about complex rollback and error handling logic. Factories that are not built on top of transacted subsystems will need to support compensation. Factories must also support enlisting in a distributed transaction.

Factories will typically maintain lots of bookkeeping information related to instantiation. This bookkeeping information must remain consistent with the SDM runtime in order to guarantee proper recovery. To facilitate this, the SDM runtime will provide a transacted storage service for component instances including factories. A well-written factory will store all its bookkeeping information in this store.

Factory port

Factories will typically expose one or more ports that can be used for component instantiation. While the port types are not mandated by the SDM

1 runtime we recommend that all component factories support the SDM_Factory
2 port. SDM_Factory is a SOAP based port that is called to instantiate new
3 component instances. The C# interface for this port is as follows:

```
4     public interface ISDMFactory
5     {
6         ComponentInstance Instantiate(
7             ComponentInstance parent,
8             Component component,
9             ComponentType componentType,
10            object args);
11
12        void Alloc(ComponentInstance allocInstance);
13
14        void Construct(ComponentInstance constructInstance);
15    }
```

16 ISDMFactory supports a three pass instantiation process:

17 Instantiation Pass: this pass will create all the component instances
18 recursively with the SDM runtime. It will not however do any
19 allocation or construction. It merely just creates the “skeleton”
20 component instances required.

21 Allocation Pass: during this pass all the relevant component factories
22 will allocate any resources needed for the instantiation.

23 Construction Pass: If the allocation succeeded then the construction pass
24 will start. This is typically the longest running pass. The factories
25 will typically do all the real work during the construction pass.

1 Factories can certainly support other port types for instantiation, but the
2 SDM runtime and Runtime APIs have a lot of helper functions that work well with
3 the SDM_Factory implementation. These APIs will certainly improve the
4 developer experience for the majority of developers.

6 Runtime-hosted component instances

8 Besides factories, the SDM Runtime will also host implementations for
9 SDM components that reference a CLR assembly using the implementation
10 SDML keyword. The referenced CLR assembly is a literal string that is the fully
11 qualified name of a CLR class. For example:

```
12     componenttype A  
13     {  
14         port pt x;  
15         implementation  
16         "MyNamespace.MyClassName,MyClrAssemblyName"
```

17 or for strongly named CLR assemblies you can specify the culture, version and
18 key:

```
19     componenttype A  
20     {  
21         port pt x;  
22         implementation "MyNamespace.MyClassName,  
23         MyClrAssemblyName, culture=neutral, version=1.0.0.1234,  
24         PublicKeyToken=9a33f27632997fcc"  
25     }
```

26 For such components, the SDM Runtime will act as the factory and it will host and
27 manage these CLR classes. This also ends the infinite recursion of factories

1 mentioned above since the base level factories are implemented as CLR
2 assemblies hosted by the SDM runtime.

3 The CLR assembly will be hosted using Microsoft's IIS Server. The
4 implementation keyword references a class that must inherit from
5 MarshalByRefObject and must implement the IRuntimeHostedImplementation
6 and the ISDMFactory interfaces. For convenience, the base class
7 SdmComponentInstance provides a default implementation for these interfaces.
8 The following is an example of a runtime-hosted CLR implementation for
9 component type A above.

```
10     public class A : SdmComponentInstance  
11     {  
12         protected override void OnCreate(object args)  
13         {  
14             // do something  
15         }  
16     }
```

16 class A is a C# class that inherits from SdmComponentInstance and therefore can
17 be hosted by the SDM Runtime. The CLR assembly for this class must also be
18 placed in the \bin subdirectory of the SDU in order for it to work properly.

19 When an instance of component of type A is created the runtime is responsible for
20 finding an available host IIS machine and instantiating the CLR code on that
21 machine. The CLR code is hosted as a .NET remoting application hosted by IIS.
22 All CLR assemblies within an SDU will share an IIS process and have their own
23 AppDomain within that process.

24 Once the CLR assembly is loaded the runtime will perform a .NET
25 remoting call to the well-defined entryptpoint on the

1 IRuntimeHostedImplementation interface. At this point the CLR class is
2 equivalent to a Component Factory and the ISDMFactory interface is consumed as
3 we saw in the previous section.
4

5 Ports and Wires

6
7 Ports and wires are the basis for communication within the SDM Runtime.
8 Ports and wires solve a number of problems that are common in service
9 deployment today:

10 Hard coding of communication information – many services typically
11 hard code the name of their server or ip addresses within their code.
12 For example, front end servers will typically hard code the SQL
13 server machine name as well as the connection information such as
14 database name, login and password.

15 Defining a communication topology – most service deployments
16 typically use DMZs as the only mechanism for defining boundaries
17 for communication. Other constraints are not enforced, for example
18 if the front end server ever needed to communicate with other front
19 end servers, this is not captured anywhere.
20

21 Discovery – finding out about new components that are added and
22 removed from a service is a typical problem faced by services today.

23 The SDM solves these problems with ports and wires. Ports are typed entities that
24 are exposed on components. A port is analogous to service access point – it is
25

1 where the component exposes well defined functionality. For example, a “storage”
2 component would define a port of type SMB.Server that can be used for filesystem
3 operations. Wires define the permissible bindings between the ports. They form a
4 communication topology that can constrain the communication paths.

5 Let’s reexamine the MyService example from above:

```
6     componenttype MyService
7     {
8         component MyFrontEnd fe;
9         component MyBackEnd be;
10
11         port http = fe.http;
12
13         wire SqlTds tds
14         {
15             fe.catalog;
16             be.sqlServer;
17         }
18
19         implementation "MyService, MyClrAssembly";
20     }
```

15 MyService contains a single wire called tds. Wires, just like components, can have
16 instances. For example, the following are two component instance of MyService
17 ms[1] and ms[2] with two different wire instance topologies.

```
18     component instance ms[1]
19     {
20         wire instance tds[1]
21         {
22             fe[1].catalog
23             fe[2].catalog
24             be[1].SqlServer;
25
26         }
27
28     }
29
30     component instance ms[2]
31     {
32         wire instance tds[1]
33         {
34             fe[1].catalog
35             be[1].SqlServer;
36         }
37     }
```

wire instance tds[2]

fe[2].catalog

be[1].SqlServer;

ms[1] has a single wire instance tds[1] that contains three port instances. ms[2] has two wire instances tds[1] and tds[2] that have two port instances each. In the first case, fe[1] and fe[2] can see each other. In the second case fe[1] and fe[2] will not see each other.

Wire instance form a physical communication topology. Port instances are members of a wire instance. They can:

- 1) Query or discover each other – the runtime API supports functions for querying and discovering other port instances on the same wire instance. All members are visible within the same wire instance. In addition, the owner of the wire instance can query the members at any time.
- 2) Receive events – members of a wire will receive events triggered by SDM operations on member port instance. See “Events” below for more details.
- 3) Constrain communication – wire instances constrain the allowable communication paths between component instances.

Port binding information

Ports are the typed entities that are exposed by a component. A port can have exactly one instance. A port instance can carry binding information which is typically everything required to establish a communication channel between

1 components. For example, the “be[1].SqlServer” port instance from above could
2 have the following binding information for connecting to the SQL backend:

3 “server=mySQLServer;uid=myLogin;pwd=myPwd;”

4 This string can be passed to ADO or OLEDB and a TDS connection can be
5 established to the backend SQL Server. The SDM runtime does not get in the way
6 of the communicating parties. It merely acts as the holder of any information
7 required to start the communication.

8 9 Port visibility and wire instances

11 Port instances on a component instance are only visible to other component
12 instance if they have been attached to the same wire instance. This is a pretty
13 powerful mechanism for building logical network topologies for services.
14 The SDM runtime also support means for automatically creating physical Virtual
15 Networks and employing packet filtering as needed in order to implement the wire
16 instance constraint. See the “Networking Architecture” document for more
17 information.

18 19 Events

21 The SDM Runtime raises certain intrinsic events as a result of operations
22 on the SDM instance space. For example, events are raised when a component
23 instance creates a port instance. Depending on the specific event, the destination is
24 either a compound component instance or the port instances on a given wire.

1 All events are delivered to the component instance on the runtime port. The
2 SDM runtime library is responsible for trapping these events and translating them
3 into a language-specific call. For example, the CLR-based SDM runtime library
4 will raise a CLR event.

5 6 Component instance events

7
8 These events are raised when a new component instance is created, or an
9 existing component instance is deleted. The destination of the events is always the
10 parent compound component instance. The events are sent to the direct parent
11 component instance only – they are not propagated up the instance tree. From our
12 example above, assume that component instance “u[1].foo[2]” asked the runtime
13 to create a new instance of the member component “c”. See Fig. 36.

14 The code for component instance “u[1].foo[2]” is currently running on
15 machine1. Using the SDM RTL it asks the runtime to create a new instance of
16 component “c”. The runtime knows the identity of the calling component instance
17 and can disambiguate and scope the operation. The new component instance is
18 created and an event raised and delivered back to the calling component instance.
19 When an instance is destroyed or fails the runtime will send the appropriate events
20 to the parent component instance and the appropriate component factories.

21 22 Port instance events

23
24 When a component instance creates a port instance or deletes an existing
25 port instance the parent component instance is notified of the changes. See Fig. 37.

1 If a port instance is attached to a wire instance all members of the wire instance
2 will be notified of the change as well as the parent component instance. This is
3 described in the next section.
4

5 Port states

6
7 Every port instance can be in one of the following states:

- 8 ■ Created – this is the state of the port when it is first created. This
9 triggers an event that is sent to the parent component instance.
- 10 ■ Attached – the port goes into this state when it is attached to a wire
11 instance. This triggers an event that is sent to the parent component
12 instance and all members of the wire instance.
- 13 ■ Online – the port goes into this state when it is ready for operation.
14 This triggers an event that is sent to the parent component instance
15 and all members of the wire instance.
- 16 ■ Offline – the port goes into this state when it wants to stop normal
17 operation. This triggers an event that is sent to the parent component
18 instance and all members of the wire instance.
- 19 ■ Detached – the port goes into this state when it is detached from a
20 wire instance. This triggers an event that is sent to the parent
21 component instance and all members of the wire instance.
22
23
24
25

- Deleted – the port is in this state when it is removed from the instance space. This triggers an event that is sent to the parent component instance.

Wire instance events

Wire instance events are raised when a wire instance is created or deleted. The destination of these events is always the parent component instance that owns the wire. See Fig. 38.

Wire instances can also contain port references to its members. This wire membership determines the destination of certain member port events. Let us continue our example from above. Assume that “foo[2].c[2]” has created a number of new instances as follows:

```
component instance universal[1]
  component instance foo[2]
    component instance c[2]
      port instance y[1]
      component instance b1[1]
        port instance x[1]
      component instance b2[1]
        port instance x[1]
      wire instance p[1]
        b1[1].x[1]
        b2[1].x[1]
```

Note that wire instance “p[1]” contains references to two port instances “b1[1].x[1]” and “b2[1].x[1]”. Let us assume that component instance “b1[1]” and “b2[2]” each run on separate machines. Fig. 39 shows the events raised when “b2[1]” changes its port state to offline.

1 Note that the “b2[1]” is hosted on Machine3 and it invokes the “set port
2 state” operation on the runtime. The runtime records the change and sends three
3 events – one to the wire instance owner “u[1].foo[2].c[2]” and two to the wire port
4 instance members “b1[1].x[1]” and “b2[1].x[1]”.

5 6 Event Delivery and Queues

7
8 The runtime will guarantee in-order delivery of events but it will not
9 guarantee a complete virtual synchrony between all members of a given wire
10 instance. In other words the SDM Runtime will allow forward progress to be made
11 even if a component instance is running slow or is dead.

12 SDM events are queued for each component instance. The operation that
13 triggered the event is considered successful if the event is successfully queued on
14 the target’s queues. The queue’s are circular in nature and can wrap around if a
15 component is severely lagging or is dead. Wrapping around will generate a new
16 “wrap-around” event. This event is sent to the component instance itself as well as
17 the parent and any owning factories.

18 19 Runtime partitioning

20
21 In order to support a large number of clients the runtime can be partitioned.
22 Due to the strict hierarchy of the SDM instance space this problem is fairly
23 tractable. The SDM runtime can be hosted on many machines across a specific
24 deployment. Each SDM Runtime instance is responsible for tracking a portion of
25 the instance space. Component instances communicate with the appropriate

1 runtime using the SDM Runtime Library. Fig. 40 shows a partitioned runtime and
2 some clients.

3 Machine 1 contains two component instances and an SDM Runtime library.
4 Machine 2 contains a single component instance and a runtime library. Machine 3
5 is hosting a dedicated SDM Runtime. Machine 4 has an SDM runtime and a
6 component instance. Note also that the two SDM Runtimes on machine 3 and 4
7 are communicating.

8 9 Partitioning

10
11 The runtime leverages the natural hierarchy inherent in the SDM to
12 partition itself. The act of partitioning involves distributing portions of the SDM
13 type and instance space across different running runtime instances. Partitioning is
14 a must for scalability. Partitioning happens differently for types and instances:

- 15 ■ Type and member space: A given runtime can contain many type
16 definitions that are typically organized within a namespace. Each
17 runtime will only need to know about the types and members that
18 are defined by the instances that it's tracking. These can appear on
19 multiple runtimes. In other words, overlap is permitted in the type
20 and member space.
 - 21 ■ Instance space: A given runtime will only be tracking a portion of
22 the instance space. The instance space is partitioned on compound
23 component instance boundaries. Overlap in the instance space is not
24 permitted.
- 25

1 This is best explained by an example; consider the following component type
2 definition:

```
3     componenttype B {  
4         port X x;  
5     }  
6     componenttype C {  
7         port Y y;  
8         component B b1;  
9         component B b2;  
10        wire P p { b1.x; b2.x; }  
11    }  
12    componenttype A {  
13        port internal Z z;  
14        component C c;  
15        wire W w { z; c.y }  
16    }  
17    componenttype universal u {  
18        component A foo;  
19        component A bar;  
20    }
```

17 This definition contains three component types A, B, and C. A is member of
18 the root universal component. B and C are members of A. It is convenient for us to
19 represent member space pictorially as shown in Fig. 41. We will use boxes to
20 represent compound components. Note that compound component members that
21 are not other compound components are described within the component box. In
22 this example, wire “w” is a member of compound component “foo” and “bar” and
23 is therefore represented within the “a” box.

24 In the instance space, there may be many instances of each component, port
25 and wire. We represent the instance hierarchy as shown in Fig. 42. The boxes here

1 represent the instance state tracked for a component instance – it is not the
2 component instance implementation code.

3 Let us assume that we wanted to partition this SDM model between three
4 runtimes – runtime1, runtime2 and runtime3. Fig. 43 is an example of partitioning
5 the instance space. In this example, Runtime1 is tracking “universal[1]”, “foo[1]”,
6 “foo[2]” and “bar[1]”. Runtime2 is tracking “foo[1].c[1]”, “foo[1].c[2]”, and
7 “foo[2].c[1]”. Runtime3 is tracking “bar[1].c[1]”. In addition, the runtimes must
8 know about all the types for the instances that it is tracking. In this example,
9 Runtime3 must know about component type “C”, “B” and “A” due to its parent
10 “bar”. It must also about port type “Y” and wire “P”.

11 The different runtimes must also maintain a relationship between
12 themselves. This relationship is mandated by the SDM hierarchy. In the previous
13 example, Runtime1 and Runtime2 must know about each in order to manage the
14 “foo[1].c[1]”, “foo[1].c[2]” and “foo[2].c[1]” relationship. Similarly Runtime1
15 and Runtime3 must coordinate work surrounding “bar[1].c[1]”. Note that
16 Runtime2 and Runtime3 do not know about each other.

17 18 Partitioning strategy

19
20 The runtime will contain enough logic to self-partition itself. The specific
21 partitioning strategy will be based on performance, capacity and SDM defined
22 constraints. This partitioning is dynamic and will change as the SDM model
23 grows.
24
25

Single-root Runtime

Runtimes that are tracking compound component instances that are all instances of a single root component instance are referred to as single-root runtimes. In the example above, Runtime1 and Runtime3 are single-root runtimes. Runtime1 has a root instance tree starting at “universal[1]” and Runtime3 has a root instance tree starting at “bar[1].c[1]”.

Multi-root Runtime

Runtimes that are tracking compound instances that do not have a root compound component instance are referred to as multi-root runtimes. In the example above, Runtime2 is a multi-root runtime since its tracking “foo[1].c[1]”, “foo[1].c[2]” and “foo[2].c[1]” which are all roots.

Service Installation

Before a service can be instantiated on a given SDM Runtime it must first be installed. The installation process involves the following steps:

Copying the Service Deployment Unit to a runtime deployment share

Calling the SDM Runtime API to start the installation

Service Deployment Units

The SDU is a unit of service deployment. It is comprised of:

SDM assembly – this is the type information for the new service. It includes all the component type, wire types and port types for that service. This assembly is a result of compiling the service SDML.

Runtime-hosted component instance code – any CLR code that is hosted by the runtime and referenced by the implementation keyword in SDML must be included in the SDU.

Other service binaries – all other binaries such as configuration files, DLLs, GIFs, HTML, SQL Scripts etc. are also considered as part of the deployment unit.

The SDU is immutable – changes to the SDU are not permitted. Once an SDU is installed it can not be changed. Certainly, one can install a new version of the SDU that upgrades and potentially obsoletes the old version(s).

SDU format

The SDU is a directory of binaries that are consumed by the SDM Runtime and potentially component factories. The directory is pretty much free form but the following structure is expected:

```
\sduroot
  \<assembly name>.<version>
  \<assembly name>.sdmassembly
    \bin
      \<Runtime hosted CLR assembly_1>.dll
      \<Runtime hosted CLR assembly_2>.dll
      \<Runtime hosted CLR assembly_n>.dll
```


\<other files and directories>

The SDU will be packaged up as a CAB file.

Implementation

The SDM Runtime is implemented as a .NET WebService running on-top of IIS Server. The SDM Store is a reliable SQL Server database. The SDM runtime webservice is a stateless webservice. In other words any state in the SDM runtime service is transient. All durable state will be written to the store at clear transaction boundaries.

The SDM runtime service can be shutdown and restated at any point an even on different machines. If it is pointed at the same SDM store all work will resume with little or no interruption.

SDM Store

The SDM runtime utilizes a durable store for the SDMs and instances. This store is typically collocated on the same machine as the SDM runtime service but it can certainly be deployed differently. The SDM store is a SQL server database that contains information about all SDM models and their instances.

This reliability and availability of the SDM store is imperative. One of the key design goals for the SDM is the ability to restart the system at the last know consistent state. The SDM therefore needs to be highly reliable and must survive catastrophic disaster scenarios. This is implemented in two ways:

1 The SDM Store will be replicated and a redundant hot backup will
2 always be available. This is implemented using Yukon's Redundant
3 Database Technology.

4 The SDM Store will be backed up regularly and the information will be
5 stored off site. The backup will be a self-consistent snapshot of the
6 current models, instances and any service state that was store in the
7 SDM Store.

8 9 Service Storage

10 The SDM Runtime will provide facilities for storage at a component
11 instance level. Every component instance can use the runtime API to store data in
12 the SDM store. At a minimum this store is a BLOB store although we are
13 considering semi-structure storage.

14 Service state stored in the runtime is guaranteed to be as reliable and
15 durable as the SDM runtime. It is also guaranteed to be consistent with other
16 runtime state. Certainly we are not advocating for all service state to be stored in
17 the SDM store instead we expect services to store sufficient information (in terms
18 of pointers) to their state. Upon recovery the service can retrieve the pointers to its
19 data and perform the necessary steps. See Recovery below.

SDM Runtime Security

Scenario Description

There are two basic scenarios that will define the security model for the SDM Runtime: the developer test-run scenario and the operator production deployment scenario. The common requirements for both scenarios are as follows:

- Ability to connect to target servers from the computer where the SDM Runtime is executing.
- Windows authentication using Active Directory domain accounts.
- Trusted subsystem model for accessing target server resources to perform install, update and uninstall operations.
- SDM Runtime implemented as a Windows Service and run as a trusted service account.
- A database (MSDE) configured to use Windows authentication and database roles that tracks SDM class, type and instance information.

Developer Test Run Scenario

A developer must be able to deploy a distributed application to one or more servers in a test environment. The target servers are either part of a standalone workgroup or in the same Active Directory domain. The computer from which the test run deployment is initiated must be in the same workgroup or domain as the target server(s).

1. The developer generates a Service Deployment Unit (SDU) package using Visual Studio.

1 2. The generated SDU is placed in a deployment folder on the computer
2 where the SDM Runtime service is executing.

3 3. Developer chooses a deployment action (install, update, uninstall)
4 and is prompted for Windows authentication credentials.

5 4. Developer is authenticated and mapped to a deployment role which
6 determines whether the authenticated user is authorized to perform the
7 requested deployment operation.

8 5. Developer selects which components to install, update or delete on
9 which target servers.

10 6. The SDM Runtime service connects to the selected target servers in
11 one of two-ways: if the SDM Runtime service is running as a trusted
12 service account in Active Directory, then it will connect as that account
13 on the target servers. Otherwise, the SDM Runtime service will connect
14 as the authenticated user, which may required an additional
15 authentication at the target server if impersonation is not possible.

18 Operator Production Deployment Scenario

19
20 An operator must be able to deploy a distributed application to one or more
21 servers in a data center environment. The target servers must be part of an Active
22 Directory domain or forest. The computer from which the test run deployment is
23 initiated must be in the same domain or forest as the target server(s).

24 1. The application SDU is placed in a deployment folder on the
25 computer where the SDM Runtime service is executing.

2. Operator chooses a deployment action (install, update, uninstall) and is prompted for domain credentials.
3. Operator is authenticated and mapped to a deployment role which determines whether the authenticated user is authorized to perform the requested deployment operation.
4. Operator selects which components to install, update or delete on which target servers.
5. The SDM Runtime service connects to the selected target servers as a trusted service account and performs the operations.

Feature Description

Behavioral Specification

The SDM Runtime is responsible for tracking all SDM classes, types and instances. The SDM Runtime will expose a set of SOAP interfaces for registering and operating over an SDM document for the purpose of deploying a distributed application.

The SDM Runtime is comprised of the following major components:

- Web Service with an associated runtime library,
- Windows Service,
- Database such as MSDE (or Yukon).

Fig. 44 shows the relationships between the SDM Runtime components, the deployment tool and the target servers. In Fig. 44, a user interacts with the

1 deployment tool UI or a command-line interface in order to initiate a deployment
2 action.

3 The runtime library provides a set of SOAP interfaces exposed by the Web
4 Service. The Web Service writes information into the database that the Windows
5 Service retrieves in order to perform a deployment action. The Web Service
6 authenticates the user to the SDM Runtime database using Windows
7 authentication and authorizes deployment actions based on roles that are defined
8 in the database.

9 In a production environment, the Windows Service will execute as an
10 Active Directory service account and the target servers will be configured to trust
11 the domain service account for administrative purposes. The Windows Service will
12 use WMI to remote to the target servers using impersonation of the service
13 account (not the user). This trusted service model should be more scalable and it
14 will minimize the need to manage target server ACLs on a per user account basis.
15 Operators will not have to be administrators on the target servers in order to
16 execute deployment operations.

17 In a test run environment, the Windows Service will execute as either an
18 Active Directory service account or as a non-privileged NetworkService account
19 in the absence of Active Directory. The latter will require impersonation of an
20 authenticated user account on the target servers.

21 UI Description

22 There is no UI for the SDM Runtime itself. The SDM Runtime will expose
23 a set of APIs which can be invoked through a deployment tool UI or through a set
24
25

1 of command-line tools. The deployment tool UI will be specified in a separate
2 document.

3 Security Model

4 The security model for the SDM Runtime is that of a trusted subsystem that
5 uses a fixed identity to access the target servers to which distributed components
6 will be deployed. The security context of the authenticated user does not flow
7 through to the target servers in this model. The basic assumption of this security
8 model is that the target servers trust the fixed identity of the SDM Runtime service
9 thereby eliminating the need to manage administrative rights for individual users
10 on the target servers. Fig. 45 shows the fixed identity trust relationship.

11 With the trusted subsystem model it is certainly possible to run the SDM
12 Runtime service under a trusted domain account or even to run it as a local non-
13 privileged NetworkService account. The key point to understand is that the
14 authorization for any deployment action is managed by the SDM Runtime using
15 role-based authorization, and that only the SDM Runtime service can perform
16 install, update and uninstall actions on the target servers once the user has been
17 authenticated and mapped to a role that permits the requested deployment
18 operation.

19 Authentication

20
21 Authentication is the process of verifying a user's identity based on a
22 credential secret known only to the user and the underlying security infrastructure.
23 For the purpose of distributed application deployment, the user will be
24 authenticated using Windows authentication either through Active Directory
25 domain accounts or local accounts. If local accounts are used, the local account

1 names and passwords on the deployment computer must be the same on the target
2 servers.

3 4 Authorization

5
6 Once the user is authenticated, authorization for performing a deployment
7 operation such as install, update or uninstall will be granted based on the database
8 role the authenticated user is a member of. Because Windows user and group
9 accounts can be members of SQL Server database roles, the basic authorization
10 sequence is as follows:

- 11 1. Web Service authenticates user using Windows authentication.
- 12 2. Web Service connects to database as the authenticated user.
- 13 3. User is mapped to a database role based on user or group account
14 membership.
- 15 4. Web Service writes deployment action information to appropriate
16 database table that can be read asynchronously by the Windows Service
17 component of the SDM Runtime
18

19
20 Notice that there is no need to manage passwords outside of the operating system
21 infrastructure nor to manage per user ACLs on the target servers.

Impersonation

Impersonation is the ability to execute code in the security context of a different account than the current process owner. Remote connections to target servers will be established using WMI with impersonation enabled. Impersonation will be based on the trusted service identity when Active Directory is present and the security context of the authenticated user when Active Directory is not available (e.g., test run environment).

Windows Service

The Windows Service component of the SDM Runtime should be run as a service account with administrative rights on the target servers. The need for administrative rights is due to the requirements of installing software on the target servers and creating various settings for IIS, SQL and the registry.

In the absence of an Active Directory domain account, the Windows Service will impersonate a user account that is authorized to perform administrative actions on the target servers. In this case the Windows Service will run as a NetworkService account which does not require passwords and is a non-privileged user on the local computer. The Windows Service will present the local computers credentials to remote computers when connecting.

1 IIS

2 SQL Server

3
4 SQL Server can operate in two authentication modes: Windows
5 Authentication mode and Mixed mode. Because Windows Authentication mode is
6 more secure than Mixed mode, SQL Server for the SDM Runtime database will be
7 configured for Windows Authentication mode only. This will prevent the sa
8 account from being used to authenticate to the SDM Runtime database.

9 Administrative privileges for the SDM Runtime database should be controlled
10 through Windows group membership in order to leverage the Active Directory
11 authorization infrastructure. By creating an Active Directory group for
12 administering SQL Server and adding specific users to the group, it will be easier
13 to control access to the SDM Runtime database without having to manage
14 passwords on a specialized account.

15 In addition to the SDM Runtime database, target servers running SQL
16 Server should also use Windows Authentication mode and manage administrative
17 access through Windows group membership. The Windows group for the SDM
18 Runtime database and the Windows group for the target servers should be different
19 groups. It is a policy decision for the customer whether or not to have one or
20 several Windows groups for administering the SQL Server machines.

21 For example:

22 SDM Runtime Administrator Group

23 User A, User B

24 SQL Server Tier 1 Administrator Group

25 User C, User D

1 SQL Server Tier 2 Administrator Group

2 User C, User E

3
4 SDM Server Overview

5 Introduction

6 What is the SDM server - The SDM Server is the set of services built around the
7 SDM. There are currently two general approaches we can take on the architecture
8 of the deployment tool. Each is outlined here.

9 Distributed approach

10
11 In this approach tools that make use of the SDM runtime and deployment
12 engine are built against a runtime OM client library which in turn communicates
13 using a web service to the SDM runtime engine and a file share for placing SDU's
14 (binaries). The SDM and deployment engines share a database of SDM entities
15 and deployment jobs. Deployment tasks are performed asynchronously by the
16 deployment engine using WMI and SMB (file share) to communicate with the
17 target machines.

18
19 Simplified approach

20
21 In this approach the client, SDM object model library, SDM engine,
22 deployment engine and installer plug-ins all run in the same process so that there
23 is no service as such. The Runtime database and binaries library can be on
24 different machines. The WMI and SMB connections to target machines are
25 directly from where the client or UI is running.

User Interface and Other Clients

The user interface for the SDM server will include:

- A wizard in Visual Studio that will provide a simple method to deploy, update or remove a test instance of an application.
- Command line tools to load SDM's, SDU's and instance requests.
- A complete UI that surfaces all the functionality of the object model and additionally provides graphical tools for composing Host models and instance requests.

Runtime OM Library

The public interface to the SDM server is through this library. It is a managed code object model and using it you can:

- Manage the SDM's in the runtime. You can load SDM's into the runtime. SDM's are strongly named and immutable and are loaded a SDM at a time (i.e. you load an SDM file not individual types, classes or mappings). You can delete SDM's from the runtime and produce the XML document for an SDM in the runtime. SDM's cannot be deleted from the runtime while there are references to it from other SDM's in the runtime or from instances.
- Manage the SDU's known by the runtime.
- Find and reflect on SDM elements (from SDM loaded in the runtime). There is no API provided for authoring a new SDM (i.e.

1 this is a read only object model over the immutable elements of the
2 SDM). This includes SDM's, SDU's, identities, versions, classes,
3 types, binding/mappings and versioning policy.

- 4 • Find and reflect on instances of components, ports, wires and
5 physical placements (the hosting relations in the instance space). In
6 the instance space each instance can be identified by a GUID, a
7 stable path or an array based path. The paths are strings and can be
8 relative. These identifiers, including relative paths allows instances
9 to be found and referenced in documents such as the instance request
10 document.
- 11 • Manipulate instances including creating, changing topology,
12 upgrading, changing settings and deleting. Instance changes are
13 made within the bounds of an instance request which provides an
14 atomic unit of update so that any errors or constraint violations will
15 result in the entire request failing. Instance requests also allow for
16 instances to exist temporarily without a binding to a host, as an
17 instance must have a host when the request is committed. It also
18 allows for many operations that will affect a single component's
19 installation or settings to be performed and have the installation or
20 settings update deferred until commit so that a single update occurs
21 on the component.
22
23
24
25

- Create sequencing within instance request when creating an instance request. Sequencing allows control over ordering of installation on the components that result from and instance request.
- Find and reflect on instance requests including getting their state including all error information, and retrying the installation/update of components affected by the request.
- Load an instance request. An instance request is an XML file that represents a set of instance space operations. This document can take advantage of relative paths to be a reusable 'script' for creating or deleting application instances.
- Generate an instance request document from an instance request in the database. Such documents are somewhat portable.
- Manage security permissions to the SDM service. This includes setting credentials used to manipulate the target machines and permissions around instance operations such as who can create instances hosted on a specific host instance.
- Subscribe to events around the functions above including, instance request installation completed. The lifetime of these event subscriptions limited by the lifetime of the process that loaded the client library (i.e. these are regular CLR events).

SDM Runtime Engine

The SDM runtime engine performs the reasoning on the SDM model and the functions surfaced by the object model.

In the distributed approach the library communicates to the runtime engine as a web service with fairly coarse calls such as load SDM, create component instance and get entire SDM (for reflecting on SDM entities). This reduces round trips to the server. The format of many of the parameters for this web service is XML with the same schema for SDM files.

In some sense the web service provides all the functionality of the SDM service with the client library simply making it much simpler to use.

In the distributed approach the engine performs the checks on permissions (see security spec for details).

Installer Plug-ins

The installer plug-ins are associated with a class host relation. They are closely related to the plug-in using in visual studio that provide the design experience for the classes and produce the associated binaries in the SDU and the deployment values. They provide the following to the functions to the SDM server:

- Installation, uninstall and reinstall components on their hosts. When an instance request results in a new component instance, removal of a component instance or a change to a component that requires a reinstall, it is the installer that takes the settings for the instance, the

1 host instance, the types associated with the component and the
2 binaries associated with those types in the SDU and performs the
3 install or uninstall of the instance. At the application layer of the
4 SDM it is most common for an installer to simply require a type
5 provided base .msi to be installed on the host (with particular
6 parameters) and a second task to execute on the host that sets the
7 appropriate settings and port views.

- 8 • Updating a component instance when its settings change or when the
9 view from one of its ports changes (either due to topology changes
10 or a visible port has settings change). At the application layer of the
11 SDM it is most common for this to be a rerun of the second part of
12 install.
- 13 • Maps the ports visible on ports to settings on an installed component
14 instance. In the SDM and component instance has port instances
15 that, as a result of some wire topology, allows the port instance to
16 see the details of other port instances, usually so that it can bind to it.
17 For example, an ASP.NET web site may have a database client port
18 instance so it can be wired to a database. When correctly wired its
19 database client port is able to see a single database server port
20 instance and the settings on that server port. This information is used
21 by the ASP.NET installer to place a connection string for the server
22 in the web.config file under the name of the client port.
23
24
25

- The installers also provide code that does the constraint checking between hosts and their guests. This check is performed by the SDM engine which is not shown in the distributed approach above. Most installers are anticipated to use a common constraint language based on XML, XPath and XQuery.
- Audit settings
- Audit existence
- Audit Full
- Audit hosted instances

Mapping settings to components.

The interface

Providing a set of base mechanisms to the installers such as execute command as local system on hosts. In the future others with provide further mechanisms that require only a net address and an account.

Interface is managed code.

Design

The following sections address how to design data centers and distributed applications that are hosted as such data centers. The designer employs the SDM to model various building blocks used in architecting the physical resources

1 employed at the data center (e.g., hardware, network, host servers) and the
2 applications.

3 4 Data Center Description

5 This section describes how to model data center components without
6 representing specific resources, such as numbers of machines. It provides a scale-
7 invariant model of the physical data center environment using the service
8 definition model (SDM) semantics.

9 A virtual data center (VDC) is a logical representation of a physical data
10 center environment that simplifies the developer's view of the data center. Ideally,
11 an IT professional or architect should be able to describe the data center in the
12 same scale-invariant manner that a developer can describe a distributed application
13 / service. The VDC is an abstraction of the server, network and storage resources
14 within the data center and their topological relationships.

15 A typical data center diagram is quite complex with multiple interconnected
16 servers, network equipment, IP addresses, VLANs, operating systems, storage, etc.
17 all expressed on a single diagram drawn using Visio or a similar tool. In addition
18 to the diagram, there are usually long documents that prescribe exactly how the
19 data center is partitioned, configured and managed.

20 An example of this complexity is the Microsoft Systems Architecture
21 (MSA) Enterprise Data Center (EDC). It should be obvious that keeping the
22 manually drawn diagrams and documents current with the state of the data center
23 over time as updates and upgrades are applied becomes a costly if not impossible
24 task. Likewise, the ability to validate the environment against the document
25 prescriptions is difficult and prone to human error.

1 The ability to represent a complex data center such as the MSA EDC in a
2 scale-invariant manner would be immensely powerful to both the developer and
3 the IT professional. The ability to describe a data center using components, ports
4 and wires provides a powerful framework within which to model and validate
5 deployment requirements that is missing in today's design and deployment
6 process.

7 One aspect of the data center description is the ability to virtualize hardware
8 and configure mechanisms for aggregated computing environments. In a
9 traditional data center environment, operators typically build out a hardware
10 environment specific to a particular application. For example, when deploying a
11 new email system into the data center, the operators will buy a set of servers, add
12 network adapters for different networks like backup and data zones, and add
13 network hardware like switches and load balancers. The deployment of the
14 hardware for an application requires extensive physical effort.

15 Not only are these manually constructed, application specific hardware
16 configurations expensive to create, but they are not easily modified; their static
17 nature results in poor resource utilization as resources can easily be moved to new
18 applications as work loads change.

19 This disclosure describes a way to create a data center virtualization
20 environment which allows operators to run a single pool of physical resources that
21 include servers, storage, and network devices. From that single pool, resources are
22 allocated and configured on demand to meet application needs. A set of resource
23 providers track the ownership of resources and know how to configure resources
24 to meet application needs.
25

1 When deploying a new application into the data center environment,
2 operators create abstract description of the resources needed by the application. A
3 request is to the services platform asking that the abstract description be resolved
4 into real resources. The services platform works with the resource managers to
5 locate resources that can fulfill the request, selects the resources which most
6 economically fulfill the request, marks the resources as used, configures the
7 resources to fit the request requirements, and places the concrete description of the
8 allocated resources into the abstract description. As the application's needs
9 change, the operator updates the resource description and asks the service platform
10 to resolve the update application description. Individual resource providers can
11 use hardware or OS specific software drivers that configure physical resources to
12 meet application needs.

13 Concepts associated with data center description include (1) a graph
14 language for describe desired resources, resource requests, and granted resources;
15 (2) a set of domain specific resource providers with knowledge of available
16 resources of a given type and the ability to configure those resources to meet
17 application requirements; and (3) a resource manager which processes resource
18 requests, communicates with resource providers to find appropriate available
19 resources, optional optimizes the choice of specific resources, asks the resource
20 providers to configure the chosen resources, and updates the resource request to
21 reflect the chosen resources.

22 23 Application Description

24 Applications can likewise be defined using the SDM semantics. This is
25 described above in more detail with reference to the SDM sections beginning in

1 paragraph 0. Fig. 20 shows a graphical user interface (UI) that allows the architect
2 to describe a large-scale distributed application in terms of SDM semantics.

4 Logical Placement of Application onto Physical System

5 Once the applications and virtual data centers are architected using SDM
6 semantics, the architect can logically try different logical placements of the
7 application elements onto the virtual hardware elements. There can be different
8 logical placements for different deployment environments (development, test,
9 production, etc.). Logical placement can be done at design time, and requirements
10 and constraints are checked and the developer is alerted of any errors or warnings.
11 The result of the logical placement is captured in a separate file, with constraint
12 checking being implemented using XPath and the XSD specified on each
13 component, port and wire class. This is illustrated in Fig. 21. The designer may
14 utilize a UI (user interface) for intuitive gestures when placing different
15 application elements onto the physical elements.

17 Design Time Validation

18 The following section addresses an approach to design time validation of
19 the logical placement of the applications onto the physical resources.
20 Enhancements to the SDM components, ports and wires add layers and mappings
21 between layers to achieve design-time validation of distributed application design
22 and deployment requirements.

23 While components, ports and wires are powerful abstractions when
24 combined with hosts, factories, resource managers and the SDM runtime, they are
25 not sufficient to deploy and manage a distributed application / service. In order to

1 create and manage the physical instances of these logical abstractions, some
2 additional constructs are involved. Those additional constructs are layers and
3 mappings.

4 5 Layers

6 Fig. 11 shows the layer abstractions defined by the SDM.

7 The application layer describes the distributable components, their
8 deployment requirements and constraints, and their communication relationships
9 in the context of an application / service.

10 The deployment layer describes the configuration and policy settings and
11 constraints for hosts such as IIS, CLR and SQL, among others.

12 The Virtual Data Center (VDC) layer describes the data center environment
13 settings and constraints from the operating system through the network topology
14 down to the servers, network and storage devices.

15 The hardware layer is describes the physical data center environment and is
16 either discovered or specified in a declarative manner using XML, for example.
17 This layer is not scale-invariant and therefore not modeled in the SDM, but is
18 included for completeness.

19 20 Mappings

21 Because the SDM is layered, there needs to be a way to bind between the
22 various layers. A mapping is essentially a binding of a component or port at one
23 layer to a component or port at the next layer down. A mapping can be described
24 as follows:

$$25 \quad M_T = [T_n \rightarrow T_{n-1}] + [T_{n-1} \rightarrow T_{n-2}] + [T_{n-2} \rightarrow T_{n-3}] [\dots]$$

1 where M represents a mapping and T represents a component, port
2 or wire and n represents the layer. The arrow symbol represents the
3 direction of the mapping which is always from a higher layer to a
4 lower layer.

5 For example, in Fig. 12 the component at the application layer named
6 MyFrontEnd is mapped to a component at the deployment layer called IIS.
7 Likewise the component named MyBackEnd is mapped to the SQL component at
8 the deployment layer.

9 10 Design-time Validation

11 The binding between a component and its host component at the layer
12 below can surface problems to the developer before the application / service is
13 actually deployed in the live data center. These problems can be due to
14 incompatible types, configuration conflicts, mismatched operations, missing
15 topological relationships, etc. For example, the attempted mapping depicted in Fig.
16 13 would result in an error because there is no potential communication
17 relationship between the IIS and SQL components at the deployment layer.

18 While the mapping from the MyBackEnd component to the SQL host
19 component could have been a valid binding based on the component and host type
20 compatibility and the lack of configuration conflicts, it is invalid because the
21 MyService SDM defined a topological relationship between MyFrontEnd and
22 MyBackEnd that does not exist at the specified deployment layer.

23 24 25 Layered Architecture

1 Fig. 48 shows a platform architecture for automating design, deployment,
2 and management of distributed applications on a distributed computing system.
3 The architecture shows multiple layers atop a base layer 302 that represents the
4 physical computer resources of the distributed computing system. An automated
5 deployment services layer 304 provides tools to convert machines into servers
6 used in the distributed computing system. Such tools allow creation, editing, and
7 deployment of OS (operating system) images. The remote programming of the
8 machine is accomplished using fully programmatic interfaces, such as WMI
9 (Windows Management Instrumentation), which is a programming interface (API)
10 in Microsoft's Windows® operating systems that allows system and network
11 devices to be configured and managed.

12 A network management layer 306 sits atop the automated deployment
13 services layer 304. The network management layer 306 allows for network
14 management and virtual topology generation. In part, the network management
15 layer supports a driver model for network computers that facilitates connection of
16 individual computers to one or more VLANs via a single physical network
17 interface connected to an associated port of the network switches. According to
18 the driver model, a VLAN driver is installed at the server and used to create virtual
19 network interfaces (VNICs) above the single physical network interface. The
20 VLAN driver creates one virtual network interface (VNIC) for each VLAN. The
21 VNICs reside just above the network interface (NIC) in the IP stack at the server
22 so that the server can handle packets passed over more than one VLAN, even
23 though all packets physically travel through the same physical NIC.

24 The driver model supports VLAN tagging to allow data packets being
25 passed over the distributed computing system to be tagged with identities of the

1 VLAN to which they belong. The network switches enforce the tagging and only
2 accept packets with tags identifying the VLANs to which the switches belong. In
3 one implementation, the network switches have both tagged ports and non-tagged
4 ports. Tagged ports of a switch are tagged with VLANs identifiers and used for
5 connection to tagged ports of other switches. This allows rapid transfer of packets
6 through the network of switches. Untagged ports of a switch are used for
7 connection to the servers or computers. When packets reach their destination
8 server, VLAN tags are stripped from the packets prior to communicating the
9 packets upstream to the servers so that the servers need not know anything about
10 the tagging.

11 A physical resource management layer 308 resides atop the network
12 management layer 306. The physical resource management layer 308 maintains a
13 physical model of the distributed computing system, tracking ownership and
14 coordinating allocation of all physical computing resources. The physical
15 management layer 308 further supports batched resource allocation, thereby
16 enabling dynamic configuration and management of physical computing
17 resources.

18 A logical resource management layer 310 sits atop the physical resource
19 management layer 308. The logical resource management layer 310 facilitates
20 allocation of logical resources requested by the distributed application. For
21 instance, the application might call for such resources as databases, load balancing
22 services, firewall, web services, and so forth. The logical resource management
23 layer 310 exposes such logical resources.

24 The next layer is the service definition model and runtime layer 312, which
25 allows description of the distributed application and tracking of its operation. The

1 service definition model (SDM) provides a namespace and context for describing
2 operations processes and an API for application introspection and control of
3 application resources. It further enables operators and developers to share
4 common application views.

5 The sixth layer atop the computing resources layer is the components layer
6 314. This layer permits definition of reusable building blocks of a distributed
7 application, which use the SDM APIs for context, naming, and binding.

8 The top layer is the operations logic layer 316, which accommodates the
9 operational aspects of the distributed application. The operations logic is
10 responsible for starting a service, growing and shrinking the service, upgrades and
11 downgrades, fault detection and recovery, and state partitioning. The operations
12 logic enables reuse of proven operational practices across deployments and
13 applications. Through use of the SDM layer, the operations logic has context to
14 better understand issues that may arise. For instance, when a failure occurs, the
15 operations logic can determine that the failure occurred at the front-end of an
16 email service, rather than just at some server in the middle of the room.

17 18 Deployment

19 The following sections address the deployment of the data centers and
20 distributed applications. It involves instantiation of the logical models, physical
21 placement of the application, and deployment of the application and data center.
22 Fig. 23 generally illustrates the deployment phase.

23 24 Instantiation

25

1 Because SDM types are scale invariant and can be created to any scale, one
2 aspect of deployment is to define the number of instances to be created for a given
3 logical component and wiring topology to physically implement the
4 hardware/application. An instance request document is created to provide a
5 declarative definition of the instances that need to be created.

6 7 Physical Placement of Application

8 Physical placement is the act of picking the specific host instance that is the
9 target of deployment. Physical placement is constrained by the logical placement
10 and constraints are revalidated during physical placement. The physical
11 placements are saved in a physical placement file.

12 13 Data Center and Application Deployment

14 The SDU, logical placement file, instance request, and physical placement
15 file are fed into the SDM Runtime. The SDM Runtime invokes the appropriate
16 installer (based on the class and host relationship), which is responsible for
17 creating a new instance on the host and configuring it to match the settings values
18 on the type. SDM Runtime will maintain a database of all instances, their final
19 setting values, and placement. A runtime API supports querying of the instance
20 space.

21 22 BIG Deployment Tool

23 Scenario Description

24 Feature Summary

1 The BIG Deployment Tool performs distributed SDM application
2 deployment for datacenter operators and for developers testing their applications.
3 It consumes Service Definition Model (SDM) applications, which includes the bits
4 of the application (SDU), mapping files, and a set of deployment constraints. The
5 user specifies placement of the application onto his/her servers and provides
6 deployment time settings. The tool installs or uninstalls instances against remote
7 machines and provides status to the operator. The operator can later add new
8 instances, decommission instances, and reconfigure the application's topology.

9 10 Scenarios

11 A large enterprise has a separate datacenter and developer organization.
12 The datacenter deploys, maintains, and hosts applications for end-users that serve
13 both employees and customers. The datacenter's topology changes infrequently
14 and closely matches the MSA EDC 1.5, which is not a BIG Computer.

15 The datacenter org provides the developers a scale-invariant abstraction of
16 its hosting policy, which we call a Logical Information Model (LIM). The policy
17 specifies the hosts' configuration including constraints on applications, allowable
18 settings, and basic topology.

19 The developer org codes and hot-fixes these applications to meet the end-
20 users' needs and stay within the datacenter's policy. The developer provides
21 deployment guidance by specifying the apps requirements and expected hosts.

22 The Application Operator uses the BIG Deployment Tool to deploy
23 applications in the datacenter. The Deployment Tool uses the developer guidance
24 and datacenter policy to ensure proper deployment. The Application Operator
25 later uses the tool to scale-out, reconfigure the apps topology, or uninstall.

Feature Description

Behavioral Specification

The overview of how the tool fits with Whidbey and other products is shown below. Note the SDM Runtime, LIM, SDM/SDU, and Whidbey are detailed in other specs. Fig. 49 illustrates an example usage flow for application deployment.

The key points to communication in Fig. 49 are (from left-to-right):

The developer delivers an application SDU, which includes the SDM, binaries, and SDU mappings. (We use binaries to mean the application bits and content.)

The Development & Datacenter Orgs are separate but share the same LIM.

On the machine running the Deployment Tool, there is a SDM Runtime with stores and APIs.

The Application Operator is responsible for the Datacenter Description; Deployment Descriptor; and uses the LIM, SDU, and SDU mappings.

An agent and “mini-factories” reside on target servers, which take the SDU, Deployment Descriptor, and Datacenter Descriptor as input for deployment.

The Agent uses a common mini-factory API to talk with the mini-factories.

The mini-factories in this example are SQL and IIS but may be extended for other products. These will do the work of install, configure, and uninstall.

Overview of settings and constraints

The BIG Deployment Tool consumes SDM applications. In order to understand how the tool will use settings and constraints, this section provides a

1 basic overview of settings and constraints with the SDM. For a full explanation
2 on settings, constraints, and the schema, see the related specs. In this discussion,
3 we do not differentiate if the settings/constraints are on the SDM meta-type, type,
4 or member.

5 With the SDM model, developers, Network Architects, and Application
6 Operators will have the ability to provide settings/constraints (Network Architect
7 and developers), SDU mappings (developers), and deployment-time settings
8 (Application Operators). These constraints and settings will be scoped per host
9 (i.e. IIS, SQL, BizTalk) each with its own schema, rules, and values.

10 Each host's exposed group of settings will be divided into those settable by
11 the application and those reserved by the host. We refer to the former as
12 application settings and latter as host settings. Furthermore, a host restricts the
13 application settings by specifying 'host constraints', and an application gives
14 prerequisites on host settings through 'application constraints'. Restrictions may
15 be on a setting's range of values, a specific value, or dependencies.

16 The following table summarizes settings and constraints for hosts versus
17 applications.
18
19
20
21
22
23
24
25

Table 1: Setting Definitions

Definition of Settings/Constraints	Example
Application settings – settings made by the developer about the application	Shopping app: maxWorkerThreads=8 401k app: maxWorkerThreads=4
Application constraints – prerequisites against the ‘host settings’ needed to run the app	Mode=WorkerProcessIsolationMode
Host settings – group of settings for all applications hosted on that resource	Mode=WorkerProcessIsolationMode
Host constraints – limits (exact values, range of values) against application settings	High-perf host: maxWorkerThreads<25 Best-effort hosts: maxWorkerThreads<5

The goal of the Logical Information Model (LIM) is to provide an abstracted view of the datacenter’s policy and deployment blockers. The LIM declares the division between host versus application constraints/settings; host settings; and application constraints. The policy that the LIM captures is authored by the Network Architect. This policy may be codified into a LIM file by the Network Architect, developer, or facilitated by the use of a canonical Microsoft LIM that is edited with Notepad.

The LIM is then used by developers to write applications and test against its representation of the datacenter. As part of the application, developers supply values for the application settings that the LIM permits, host constraints for where the app will run, and metadata about placement of components onto hosts. Developers provide guidance on the placement of the app onto hosts through a mappings file. Unspecified settings will be passed through as deployment-time settings that Application Operators will provide (i.e. IP address or App_pool_ID).

1 A basic example would be a Network Architect specifies different host
2 constraints for customers buying services on *High-perf* versus *Best-effort* hosts.
3 The host constraint may limit the number of IO or WorkerThreads differently.
4 The *High-perf* and *Best-effort* host settings in this example are the same, using
5 IIS_6's new mode. The developer writes two applications with different budgets
6 and requirements. The first *Shopping* application wants more WorkerThreads.
7 The *401K* application is less discriminating. Both applications constrain (require)
8 running in WorkerProcessIsolationMode. Fig. 50 illustrates application versus
9 host 'settings and constraints'.

10 11 Phases of Deployment

12 Using the BIG Deployment Tool, there are four phases around SDM
13 application deployment shown below. Fig. 51 illustrates example phases for a
14 deployment tool.

15 Initial Phase is where the LIM is produced to represent the datacenter in a
16 scale-invariant manner and then used to create a hardware classification file
17 (Datacenter Descriptor).

18 App Deployment Phase is when the developer codes against the LIM and
19 uses the Deployment Tool APIs to test and debug his/her SDM application.

20 Install Phase is where the Application Operator installs apps on an already
21 configured machine.

22 Running Phase is when the Application Operator scales-out, reconfigure the
23 topology, or uninstalls an already running app.

24 Note throughout this document and especially in the flow charts, we use the
25 term "deploy" to include all the necessary host settings/constraints check, flagging

1 host versus app incompatibilities, writing app settings, and calling the mini-factory
2 actions. Mini-factory actions are all those that perform install; uninstall;
3 configuration; and hook into Fusion, MSI, or future Microsoft installers.

4 5 Initial Phase

6 The Initial Phase is when the LIM and Datacenter Descriptor are created.

7 The datacenter's Network Architect selects and downloads the closest
8 matching, digitally-signed LIM from Microsoft.com. The Network Architect then
9 edits the file to reflect the desired datacenter policy, including network topology,
10 permitted application settings, and hosting constraints.

11 Additionally, a LIM can be authored in Visual Studio Whidbey's design
12 surface. The process flow would then be a Network Architect gives the developer
13 org all relevant policy and topology information, which today are captured in
14 Word docs and Visio diagrams. The developer then creates the appropriate LIM
15 describing the datacenter and iterates with the Network Architect to ensure
16 correctness.

17 Once the LIM is created, the datacenter org then classifies their hardware
18 according to the LIM by creating a Datacenter Descriptor file. The Datacenter
19 Descriptor maps the LIM components against running hardware, which we call the
20 act of classification. Therefore, the Datacenter Description is not scale-invariant
21 and includes machine specific details like IP address. The following figure
22 visualizes a Datacenter Descriptor but does not suggest UI. Note a LIM would
23 have the concept of "IIS gold" and "IIS silver" logical hosts. In the Datacenter
24 Descriptor, these logical hosts are mapped to physical machines, thereby we have
25 an IIS[1] gold is on IP address 192.168.11.2, IIS[2] gold is on IP address

1 192.168.11.3, etc. Fig. 52 illustrates an example visualization of a datacenter
2 description.

3 Note as the Datacenter Operator installs/configures the servers, network,
4 resources, and everything below the application, actions need to stay within the
5 LIM. (Remember the Datacenter Operator is responsible for everything below the
6 application.) Both the Network Architect and Datacenter Operator perform their
7 tasks outside of the Deployment Tool.

8 9 App Development Phase

10 In this phase, the developer codes against the LIM and uses the BIG
11 Deployment Tool APIs for test/debug deployments. This LIM was either provided
12 by the datacenter or codified by the developer org on behalf of the datacenter (as
13 described above).

14 The Deployment Tool APIs enable two scenarios for Visual Studio
15 Whidbey to perform their “F5” and “Test/debug” deployments. The F5 and
16 Test/debug deployments are to a single developer box and multiple machines
17 respectively. In the F5 scenario, the necessary bits are already on the target single
18 developer box. The Test/debug case requires that the Deployment Tool transfer
19 bits to targets machine(s), as in normal deployments. However, both F5 and
20 Test/debug scenarios enable developers to be warned of conflicting settings and
21 overwrite both application and host settings. (Normally, only the application
22 settings can be written by the Deployment Tool.) Note these VS scenarios will not
23 use the SDM Runtime. Fig. 53 depicts these VS scenarios.

24 The important caveats for the Visual Studio “F5” and “Test/debug”
25 scenarios are:

1 The BIG Deployment Tool API's will be called from VS through a wizard.

2 The VS wizard will select machines to deploy against and take deployment-
3 time settings (i.e. IP_address or App_pool_ID=17).

4 VS will implement the user interface.

5 In the F5 scenario, the SDM, SDU, binaries, and all bits are already on the
6 target single development box. Thus, writing settings is all that is needed.

7 In the Test/debug loop, "Deploy" includes writing the necessary settings.

8 Both scenarios flag when settings conflict and allow overwriting the target
9 machines' settings, including host and application.

10
11 Not shown in Fig. 53 is the developer is coding the application against the
12 LIM and the notion of SDU mappings to a LIM. (For more on the LIM, see the
13 LIM/LID specs.) The developer delivers to the Application Operator the SDU,
14 which contains the SDM, binaries, and SDU mapping files.

15 16 Install Phase

17 For the Install phase, the operator is provided the application (SDU with
18 mappings) and Datacenter Descriptor (which extends the LIM).

19 For Fig. 54 describing application installation, the following caveats are
20 important:

21 The Application Operator launches the tool (GUI/CLI).

22 Copies and Loads the application with all the files and Datacenter
23 Description.

24 The application is registered in the SDM Runtime.
25

1 Application Operator selects the host/machine of the
2 application_components. (Examples are given in the next section.)

3 During this selection (we call mapping), constraints are being checked
4 against the Runtime's view of the world. We do not make guarantees if you
5 modify the settings outside of this tool causing a disjoint view.

6 Deployment performs host versus application constraints/settings check and
7 installs. (Note the implementation may be much more complex with caching files
8 and settings ACL's on the cache to avoid network flakiness.)

9 Tool makes it clear through UI or documentation that we do not handle
10 stateful data (such as populating SQL databases).

11 Above steps generates a Deployment Description, which can be reused for
12 that specific deployment or modified. (Examples given in the next section.)

13 A "Preview" function allows the Application Operator to get a list of the
14 changes the tool will make. The user can then rerun the tool using the Preview
15 generated Deployment Descriptor.

16 An already generated Deployment Descriptor can be loaded and ran,
17 assuming the SDM Runtime knows of the application, application bits are still
18 available, and the same Datacenter Descriptor is valid. Fig. 54 illustrates an
19 example install scenario.

20 21 An example of specifying deployment

22 To clarify the flow of data needed to specify deployment, we use the
23 example of MSN constraining their datacenter through a LIM.

24 The LIM may be digitally signed, time-stamped, and versioned. The
25 development org uses the LIM to code a two tiered application hosted on MSN

1 datacenter's hosts (IIS and SQL servers). The developer specifies the host on
2 which a component should be hosted, generating a SDU mapping file. We show
3 this MSN example in Fig. 55.

4 The following are important about Fig. 55 and the flow of data in
5 application deployment:

6 The SDU includes the SDM.

7 The developer maps SDU components to the LIM (MSNdatacenter.LIM),
8 creating a SDU mappings file. The mappings are the allowable placements.

9 The Datacenter Description classifies actual/physical servers according to
10 LIM components and is not scale-invariant.

11 The SDU, SDU mappings, Datacenter Description, and user input feed into
12 the Deployment Tool to create a Deployment Descriptor.

13 The Deployment Descriptor specifies the components (from the SDU) to
14 install on which machines (from the Datacenter Description).

15 The Deployment Descriptor takes deployment-time settings such as URLs.

16 Fig. 55 illustrates an example of generating a deployment descriptor file.

17 In the example above, the SDU mappings file says that the developer binds
18 the SDM Component 2TierApp.MyWeb to the MSN constrained host
19 Component MSN9.IIS_MSN and the same for 2TierApp.MyDB →
20 MSN9.SQL_MSN. (We specify the Compound Component to disambiguate in the
21 event of multiple MyWebs.)

22 The Network Architect edits the MSNdatacenter.LIM that describes how
23 the IIS and SQL constraints and settings are configured. This LIM is scale-
24 invariant because it describes IIS and SQL hosts, not specific machines running
25 IIS or SQL. The datacenter then derives a Datacenter Descriptor, which says

1 which machines are running IIS and SQL as configured in the LIM. We use the
2 notation IIS_MSN[1] and IIS_MSN[2] to signify there are two machines
3 running the IIS_MSN component.

4 The BIG Deployment Tool takes as input the SDU, SDU mappings,
5 Datacenter Descriptor, deployment settings (provided by the user), and generates a
6 Deployment Descriptor. In our example, the Deployment Descriptor specifies a
7 deployment. Meaning running it will cause software to be installed/scaled-
8 out/reconfigured/uninstalled on target servers.

9 As detailed in the Deployment Descriptor text, an instance of MyWeb
10 (MyWeb[1]) will be installed on server IIS_MSN[1], MyWeb[2] on server
11 IS_MSN[2], and MyDB[1] on server SQL_MSN[1]. Deployment-time
12 settings are provided by the Application Operator such as IP address or
13 App_Pool_ID. Note this Deployment Descriptor may be reused provided the files
14 it depends on exist.

16 Running Phase

17 Scale-out [in] scenario

18 For an already running application, the scale-out [in] scenario allows the
19 Application Operator to add [delete] a component, port, or wire. An example of
20 the usefulness of this feature would be the Joe_Millionaire website experiences
21 dramatic increase in traffic and wants to scale-out for just the regular TV season
22 and scale-in afterwards (or nightly).

23 In the flow chart for application scale-out [in], the following are the
24 important points:

25 Scale-out [in] is a subset of installing.

1 The Application Operator selects a running SDM application and can:

2 Add components, ports, wires, and enters deployment-settings.

3 Delete components, ports, and wires.

4 Scenario can be run from a previously generated or modified Deployment
5 Descriptor. (Provided the earlier caveats of having the same Datacenter
6 Descriptor/LIM, access to application, and SDM Runtime still has the app
7 registered.) Fig. 56 illustrates an example scale-out scenario.

8 9 Topology-reconfiguration scenario

10 The topology-reconfiguration allows the Application Operator to rewire a
11 running app without uninstalling, reinstalling. Examples of rewiring would be
12 changing your front-end databases to now point to a new back-end database.

13 The important points in the topology-reconfiguration are:

14 This scenario differs from scale-out in that it allows editing of an existing
15 port and wires without uninstall, reinstall.

16 It potentially allows users to “bridge” two different SDM applications.

17 Fig. 57 illustrates an example topology-reconfiguration scenario.

18 Topology-reconfiguration is useful in failure cases where you do not want
19 to redeploy the entire application. As an example, Passport stores all my credit-
20 card-numbers on a backend and made available through an IIS front-end. The
21 front-end fails and I do not want to redeploy/migrate data. Instead, I deploy a new
22 front-end (as part of normal install) and rewire the new front-end to my Passport
23 database.

24 An example of the bridging with topology-reconfiguration would be if the
25 beta_MSN10 app wanted to share MSN9 app’s databases. The Application

1 Operator deploys the beta_MSN10 normally. Now, the beta_MSN10's front-ends
2 need to talk to MSN9's database, requiring a reconfigure (and new wire) on
3 MSN9's database.

4 5 Uninstall Scenario

6 With the uninstall scenario, the Application Operator selects the application
7 and all running instances are deleted and the Runtime is updated. The user does
8 not select the exact instance to uninstall because that is possible through the scale-
9 in scenarios.

10 The following points are important for the uninstall scenario:

11 Uninstall can be performed through an existing (potentially edited)
12 Deployment Descriptor.

13 The user selects the application to uninstall and all instances are removed.

14 Stateful content must be destroyed outside of this tool through existing
15 means.

16 Fig. 58 illustrates an example uninstall scenario.

17 18 Management

19 The following sections address the management of the data centers and
20 distributed applications after they are deployed. A model-based management tool
21 is first described, followed by discussion of an introspection/tracking mechanism
22 and operational logic.

23 Model-based Management

24 Model-based management (or Ops Logic) is processing which will accept
25 event triggers from the physical environment based on definitions of operator and

1 application developer intent and policy in an SDM-based model of the application
2 and will active and orchestrate a series of tasks or processing with in the context of
3 the model, which will instigate change and will provide consistency between the
4 model and physical world.

5 A trigger or request or other threshold will be an event aimed at a particular
6 instance in the SDM. The component instance will receive the trigger and based
7 on other details about itself in the context of the overall application and hardware
8 environment represented in the SDM, it will kick-off a sequence of steps to
9 address the issue identified by the trigger. It is the context of the application and
10 the resources from the SDM which gives this automation its richness and ability to
11 provide more manageability to the Operations staff of the service.

12 Fig. 59 is an overall architecture for the model (BIG) and management
13 pieces of what we call Ops Logic or Model-based Management. To summarize the
14 proposed flow of processing in the overall architecture:

- 15 • An application developer will be able to define an aggregate model (SDM)
16 of a new application, or the classes of component types that will make up
17 the end-user application or service.
- 18 • The developer or an operations developer will be able to also add “operator
19 intent” to the model by annotating component types in the model with
20 policy and guidelines for operation, such as setting a minimum number of
21 servers that must be running.
- 22 • The SDM Run-time or unit model of instances implemented for a particular
23 implementation of the application will be held in the Unit Model. There is
24 a one-to-one correspondence between the instances holding the desired
25 state of each machine and a physical machine.

- The resource managers of BIG will work with the Unit Model to implement change in the physical world of servers.
- Each server will in part be managed by BIG and in part may be managed outside of the model by operators.
- Between the aggregate model and the unit model is one type of model-based processing to orchestrate change and implement operator intent through the model to physical machines.
- Another type of model-based processing will flow the other way and provide consistency between the physical space and the model.
- In the management area, the Monitoring System will be collecting events and grouping them into alerts.
- Components subscribing to events and alerts will be notified of important events. The event information will flow to the subscribing component with information about the run-time SDM unit or instance involved, which provides the mapping to the model.
- If an event is an operational trigger, the event will trigger the model-based processing which can instigate change in the physical machines through a sequence of orchestrated ops tasks.

Fig. 60 shows representative layers of management. This is a blow-up of the model section of the overall architecture diagram above which has been turned horizontally so that the aggregate model corresponds to the SDM and the unit model corresponds to the SDM Instance Space. The Overall Resource Manager manages requests to the individual Resource Managers (also called Factories).

Introspection/tracking mechanism

Given a trigger such a user request, a hardware trigger, or a hardware threshold being hit, an appropriate operational process will be activated. The operational process will be a set of operations tasks that will be executed. The execution of operational tasks requires processing by orchestration because each task is a transaction which may be long-live and requires initiation and completion before the next task. The engine which oversees this sequence of activity to execute operational processes is the orchestration engine for Ops Logic.

Applying orchestration to a sequence of operations tasks on potentially distributed servers or hardware resource is a unique approach. These properties of Ops Logic make a more sophisticated approach to transaction processing:

- Long-lived – Operational processes may run for long periods of time, such as days or months.
- Asynchronous – A trigger or event may start a transaction or process, but can not wait until the triggered task is complete to process other events.
- Transacted – The steps in an operational process are actions that have an agent who starts or send it, an agent who receives and processes it and a compensation process that backs-out the changes if the task should fail.
- Durable – Ops processes need to be able to last for a long time without becoming damaged or unstable.
- Highly-available – Being available as much as possible reliability is a requirement for operational processes of the highly-available BIG computer and services.

1 Ops Logic will provide operations and application developers the
2 opportunity to codify and standardize sequences of operations actions based on a
3 trigger in the BIG environment. Once a trigger is raised, the relevant sequence of
4 tasks will be activated. The steps for a particular situation may include a
5 command to an individual machine, a change in an application component
6 instance or in the model or human steps. Each step is a transaction which has a
7 start and an end and may succeed or fail. By using an orchestration engine to step
8 through these tasks, the process will be managed, tracked and reported upon. The
9 orchestration engine will initiate a task, watch its progress and note its completion
10 or failure. Orchestration will also enable alternative actions to be taken in the
11 event of partial or complete failure, depending on how the operations process has
12 been defined. See Fig. 61.

13 14 Resource Manager

15 The Resource Manager is responsible for allocating logical and physical
16 resources within the distributed computing system. The Resource Manager
17 discovers available hardware, processes resource allocation requests, and tracks
18 ownership of logical and physical resources. By providing an interface to a
19 dynamic pool of resources, the Resource Manager provides the bed-rock for
20 availability and scalability within the server.

21 The Resource Manager owns and controls all hardware in the distributed
22 computing system including both computers and network devices such as
23 switches. Access to hardware resources in the system is controlled through the
24 Resource Manager. In addition, the Resource Manager provides base mechanisms
25 for controlling logical resources such as load balancing groups.

1 The Resource Manager provides a common API for all resource
2 management within the system. Services and the runtime converse through the
3 Resource Manager API to make resource queries, allocate resources, change
4 resource requirements, and free resources.

5 6 BIG Resource Manager

7 Introduction

8 Feature Summary

9 BIG defines a distributed service runtime, a common hardware reference
10 platform, and a resource manager. The distribute service runtime provides a
11 service with a skeleton defining the service components, their interrelationships,
12 and an execution environment for scalability and availability policy in the form of
13 operations logic. The hardware reference platform defines a common hardware
14 structure that enables services to run on systems ranging from one to thousands of
15 computers.

16 The BIG Resource Manager is responsible for allocating logical and
17 physical resources within the BIG computer. The Resource Manager discovers
18 available hardware, processes resource allocation requests, and tracks ownership
19 of logical and physical resources. By providing an interface to a dynamic pool of
20 resources, the Resource Manager provides the bed-rock for availability and
21 scalability within the BIG machine.

22 This document describes the goals, architecture, and implementation of the
23 BIG Resource Manager. Chapter 1 describes goals and driving scenarios. Chapter
24 2 describes the architecture of the Resource Manager and its associated Resource
25 Providers. Chapter 3 describes implementation details and APIs.

Discussion

The BIG Resource Manager is responsible for management of allocation and usage of resources within the BIG computer. The BIG Resource Manager owns and controls all hardware in the BIG computer including both computers and network devices such as switches. Access to hardware resources in the BIG computer is controlled through the Resource Manager. In addition, the Resource Manager provides base mechanisms for controlling logical resources such as load balancing groups.

The BIG Resource Manager provides a common API for all resource management within the BIG computer. Services and the BIG runtime converse through the Resource Manager API to make resource queries, allocate resources, change resource requirements, and free resources.

Resource Providers

While the Resource Manager provides a common interface into resource management, knowledge of actual resources comes from a set of Resource Providers. A Resource Provider has specific knowledge about the existence and management of a particular class of resources. For example, the Network Resource Provider knows about the existence and specifics of managing VLANs. Other Resource Providers in BIG include a Physical Device Provider, an IIS VRoot Provider, an SQL Database Provider, a CLR AppDomain Provider, and a Win32 Surface Resource Provider.

1 Resource Providers extend the Resource Manager with resource-specific
2 knowledge. Resource Providers manage the conversion of resource specific
3 requests to a common query format. Resource Providers extend the Resource
4 Manager API with resource-specific configuration APIs through a provider helper
5 DLL. Finally, Resource Providers add appropriate state into the Resource
6 Manager data store to allow tracking of resource specific information.

7 Higher-level Resource Providers build on lower-level Resource Providers. For
8 example, the IIS VRoot Resource Provider allocates machines through the
9 Physical Device Resource Provider. Layering of Resource Providers minimizes
10 redundancy and increases uniformity of resource management.

11 A strong analogy can be drawn between the I/O management system in
12 Windows and the resource management system in BIG. Like the Windows I/O
13 Manager, the BIG Resource Manager provides a common API, common logic for
14 resource access control, a common resource tracking, and a common mechanism
15 for walking requests through a diverse set of providers. Like Windows Device
16 Drivers, BIG Resource Providers extend the management system with specific
17 knowledge for controlling a distinct class of resources. The BIG Resource
18 Manager, like the Windows I/O manager provides a model for unifying diverse
19 resources under a common umbrella.

20 21 Automatic Resource Management and Optimization

22
23 The BIG Resource Manager frees data center operators from direct
24 involvement in the allocation and placement of components on resources. For
25 example, when a new service is installed into the BIG computer, operators do not

1 need to decide on which computers to place the service. Operators only need to
2 grant the service a resource quota; the Resource Manager then decides how to
3 optimally place the service within the BIG computer in order to conserve limited
4 shared resources such as core network bandwidth.

5 The base set of trusted Resource Providers participates in the optimization
6 of component placement with the Resource Manager. Resource Providers
7 participate in placement optimization by making providing the Resource Manager
8 with placement choices and provider-specific relative cost preferences. The
9 Resource Manager then balanced global concerns with each Resource Provider's
10 local concerns to maximize efficiency and minimize resource usage.

11 Optimal component placement is an ongoing concern. Over time, the
12 resource needs of individual services shrink and grow. The available physical
13 resources change as new equipment is added to the BIG computer and older
14 equipment is decommissioned. The Resource Manager periodically re-examines
15 placement decisions and evaluates the merit of moving components. Resource
16 Providers participate in the placement reevaluation by provide the Resource
17 Manager with costs of moving components. Movement costs can range from
18 infinite for a non-movable store to quite small for a stateless IIS component.

19 The BIG Resource Manager frees operators from concern about resource
20 allocation and component placement. The Resource Manager also frees
21 developers from the need to write complex allocation logical; instead, developers
22 simply provide the Resource Manager with a graph of resource requirements. The
23 Resource Manager takes into account both local and global resource requirements
24 to optimally place components within the BIG computer.
25

1 Feature Description

2 Execution Environment

3 The BIG Resource Manager runs as a CLR managed service backed by
4 Highly-Available SQL. It is expected that each BIG machine will only have a
5 single Resource Manager replicated across the pair of SQL servers in the HA SQL
6 cluster.

7 The BIG Resource Providers execute within the BIG Resource Manager
8 process. The Resource Manager gives the Resource Providers an asynchronous
9 execution environment within which to operate and a shared database into which
10 they store their state. It is expected that all Resource Providers will be CLR
11 managed code using the BIG operations logic model.

12 All Resource Providers keep their state in the Resource Manager database.
13 Resource Providers can create their own tables as needed to meet their
14 management requirements.

15 The Resource Provider's state in the Resource Manager database is
16 authoritative. So, for example, the IIS metabase is a cache of the data in the
17 Resource Manager database. If an IIS VRoot entry is found in the IIS metabase
18 with no corresponding entry in the Resource Manager database, then the VRoot in
19 the metabase is deleted.

20 All resource allocation and de-allocation requests are unified within
21 transactions. Resource Providers that execute within exclusively within the
22 Resource Manager process using exclusively the Resource Manager database.
23 Even aggregated, cross-provider resource requests execute in deterministic, non-
24 distributed manner. This greatly simplifies the design and implementation of
25

1 Resource Providers and ensures that resources are never lost between servers in
2 failure scenarios.

3 The BIG Resource Manager separates resource allocation and resource
4 initialization into two distinct, separate acts. Resource allocation is a non-
5 distributed, deterministic operation that executes exclusively within the Resource
6 Manager process. Resource initialization on the other hand is an inherently
7 distributed and non-deterministic process.

8 Resource allocation is typically prefaced by a depth-first operations logic
9 phase in which components are instantiated, connected with wires, and attributed
10 as necessary with resource requirements.

11 By separating component instantiation and resource allocation from
12 resource initialization, the BIG Runtime and services can use common error-
13 handling mechanisms whether a resource is unavailable because it has not
14 completed initialization or it is unavailable because the device on which it resided
15 just vaporized. Resource initialization will be driven typically by a state machine
16 that saves state in an HA SQL store like either the Resource Manager database or
17 the SDM database.

18 19 Resource Providers

20
21 The BIG Resource Manager owns all resources in the BIG computer.
22 Through resource-specific Resource Providers, the Resource Manager is extended
23 with knowledge specific to distinct classes of resources. The Resource Manager
24 provides storage, management of aggregate resource operations, and acts as a host
25 for resource providers.

1 The BIG Resource Manager provides a small, specific set of resources
2 through a limited number of Resource Providers. While small in number, it is
3 expected that the basic set of Resource Providers will cover the requirements of
4 most, if not all, of the target customers. The following Resource Providers are
5 expected in the first product release:

- 6 ■ Physical Resource Provider (raw devices)
- 7
- 8 ■ Network Resource Provider (VLANs)
- 9
- 10 ■ External Resource Provider (DNS names, external IP addresses)
- 11
- 12 ■ IP Load Balancing Group Resource Provider
- 13
- 14 ■ IIS VRoot Resource Provider
- 15
- 16 ■ SQL DB Resource Provider
- 17
- 18 ■ CLR AppDomain Resource Provider
- 19
- 20 ■ Win32 Surface Resource Provider (a Win32 program)

21 Creation Pattern

22 Typically resource management will be driven by operations logic
23 packaged as CLR managed code running. The operations logic will be written to
24 the “disembodied object” pattern in which a CLR managed object represents the
25 target component. The disembodied object is responsible for allocation any
logical or physical resources needed by the component, initialization those
resources, and eventually deconstructing and releasing those resources when the
component is no longer needed.

1 A call like,

```
2 FrontEnd f = new FrontEnd(); // Instantiate the disembodied  
3 object.
```

4 results in the creation only of a disembodied object, a CLR class with an
5 component instance record in the runtime database, but nothing more.

6 Operations logic interacts with FrontEnd f to set parameters, like scaling
7 requirements, etc.

8 The disembodied object, FrontEnd f in this case, takes part in any resource
9 allocation by responding to a request for a graph of desired resources and a
10 subsequent setting of resources,

```
11 r = f.GetResourceGraph(); // Ask f to produce the logical resource  
12 request graph (recursively if f is compound).
```

```
13 rgo = BigAllocateResources(rgi); // Ask the Resource Manager to  
14 do the global allocation request.
```

```
15 f.SetResources(rgo); // Notify f of the actual resources  
16 allocated (recursively if f is compound).
```

17 The disembodied object directs all object initialization, like say formatting a disk
18 and laying down an image:

```
19 f.BeginConstruct(); // Kick off the construction/initialization state  
20 machines.
```

```
21 f.EndConstruct(); // Get results when construction has finished  
22 (this is just the .NET async pattern).
```

1 Furthermore, the disembodied object's lifetime exceeds that of the represented
2 object with the disembodied object directing destruction. The previous statement
3 does not prohibit object quiescence.

```
4         f.BeginDestruct();    // Kick off the destruction state machines.  
5         f.EndDesctruct();    // Get results when destruction has finished.
```

6
7 The disembodied object also releases its resources:

```
8         f.ReleaseResources();
```

9
10 After which it can be destroyed:

```
11  
12         f = null;
```

13
14 There are a couple of things worth noting. Because f is just a disembodied
15 object and because resource allocation is distinct from resource
16 initialization/construction, the following lines can all be placed in a single
17 deterministic transaction. It can even be a non-distributed transaction provided the
18 RM DB is in the same SQL as the SDM DB.:

```
19         BeginTransaction();  
20         FrontEnd f = new FrontEnd;  
21         r = f.GetResourceGraph();    // Ask f to produce the logical  
22         resource request graph  
23         rgo = BigAllocateResources(rgi); // Ask the Resource Manager to  
24         do the global allocation request.  
25         f.SetResources(rgo);    // Notify f of the actual resources  
        allocated.
```

1 EndTransaction();
2

3 All Resource Providers, at some point, will invoke distributed operations,
4 but not during the BigAllocateResources() call. An implementation of a given
5 Resource Provider may leverage distributed code through its own SDM modeled
6 service.
7

8 Placement Optimization

9 First, in the context of this discussion, I'd like to define the following terms
10 with respect to placement optimization:

11 I. Local Optimization: Optimization isolated to a single component
12 factory, by implication ignoring the effects on placement within other
13 component factories.

14 II. Aggregated Optimization: Optimization taking into account multiple
15 component factories. For example, optimization which considers the
16 placement of both IIS applications and SQL databases.

17 III. Global Optimization: Optimization (including movement of existing
18 components) of the entire system, i.e. all of the applications in a BIG
19 computer. Global optimization differs from aggregated optimization
20 primarily because it has the option of moving existing components.
21

22 Unless I have misunderstood people's positions, I think everyone agrees on
23 the following:
24
25

1 I. BIG V1 should provide an aggregated allocation API. The
2 aggregated allocation API takes as arguments a collection of component
3 and wire instances with configuration parameters on the component and
4 wire instances in the SAM. In a single transaction, the aggregated
5 allocation API calls into the component factories to reserve the necessary
6 resources. [Note: I have specifically used the term aggregate instead of
7 batch to highlight the fact that the allocation may include differing
8 component factories. Note that I have not said “aggregated optimized
9 allocation API” in the point.]

10 II. In the long term, BIG should provide global placement optimization.
11 The goal of global placement optimization is to re-arrange the placement of
12 component instances within the BIG machine to optimize certain
13 properties, the primary property being the usage of the BIG machine’s
14 shared resources.

15 III. Aggregated placement optimization can be occur at initial allocation
16 or can take the form of global optimization later with controlled application
17 consent. The easiest time to affect placement is when a component instance
18 is initially allocated.

19 IV. Movement of a component after initial placement can be very costly,
20 or even prohibitively expensive. Moving a large SQL backend can be
21 extremely costly and may seriously impair application availability.
22 Movement of a component should consider the wishes of an application.

23 V. In long-running applications, movement of components will be
24 inevitable even without global placement optimization. Hardware may fail
25 unexpectedly. Hardware will definitely be decommissioned due to normal

1 depreciation and life-cycle constraints. This implies that any long-running
2 application will ultimately require some mechanism for moving
3 components. Whether these mechanisms are leveraged by global
4 placement optimization is orthogonal to the existence.

5 VI. Long-running application will support migration of some form for
6 upgrades. The mechanisms for rolling upgrade, for example, might be
7 leveraged by global placement optimization. For example, if an
8 application's rolling upgrade policy is to bring a new front-end online and
9 decommission the old one, then that allocation of the new front-end is a
10 perfect time for optimizing its placement. Upgrade provides a window of
11 opportunity for global placement optimization.

12
13 Based on feedback from other in the team, I would like to propose the
14 following for BIG V1:

- 15 1) BIG V1 provides a batch allocation API. The batch API takes as
16 arguments a collection of component and wire instances with configuration
17 parameters on the component and wire instances in the SAM. In a single
18 transaction, the batch API calls into the component factories to reserve the
19 necessary resources.
- 20 2) BIG V1 formalizes the movement of components. At a minimum this
21 should include a standard component interface for taking a component
22 offline and bring it back in another location. Think of it as the component
23 equivalent of ISerialize. This formalization would be used by operation
24 logic to perform rolling upgrades and cloning of entire front ends. It might
25 also be used for partitioning SQL back ends. It would used when

1 decommissioning hardware, etc. We should have the concept of a movable
2 component, and what it means to move different types of component, how
3 to estimate the cost, etc.

- 4 3) BIG V1 provides an aggregated placement optimizer. The complexity of
5 the optimizer is tuned to meet the needs of the development cycle. It may
6 be as simple as a crude clustering optimizer or much more sophisticated.
- 7 4) The aggregated placement optimizer is used by the batch allocation API
8 during initial placement. Component factories cooperate with the
9 placement optimizer to aid its decisions.
- 10 5) Throughout application lifetime the aggregated placement optimizer may
11 be invoked periodically to move component instances to perform global
12 placement optimization. The optimizer may leverage windows of
13 opportunity presented naturally by an application. It may also ask an
14 application to consider component movement at other times. Basically, the
15 global optimization just leverages the aggregated placement optimizer and
16 the pre-existent support for movable components.
- 17 6) BIG V1 IIS application component factory implement movable
18 components, subject to application allowance. It is quite likely that much
19 of the benefits of global placement optimization can be realized by ignoring
20 heavy components such as SQL databases and moving VRoots. IIS also
21 naturally supports operations such as drain which facilitate movement of
22 VRoots. In effect, the IIS VRoot component factor becomes the V1 poster
23 child for component movement and global placement optimization.

Modeling Physical Resources

Underlying the entire resource management system is a hardware resource graph. The hardware resource graph describes the totality of hardware resources and their connectivity available to the BIG Resource Manager. The hardware resource graph includes servers, network devices, and network topology. Additionally the hardware resource graph can contain information about power grids and physical containment relationships.

The hardware resource graph consists of three basic elements: entities, connectors, and connections.

An entity is the fundamental unit of hardware accessible by software. Examples of entities include servers, disk drivers, network devices, etc.

A connector is a physical interface to an entity. A connector is always associated with exactly one entity. Examples of connectors include network interfaces, IDE interfaces, AC power connectors, and physical containership, etc.

A connection is a physical relationship between exactly two connectors. Examples of connections include network cables, IDE cables, AC cables, etc.

All three element types, entities, connectors, and connections, have associated properties. The properties are tuples consisting of a property name, maximum value, and available value.

All three element types can have duals. A dual is a peer used for fail over. An element and its dual are always allocated together to provide redundancy necessary for high availability. Typical examples of duals include fail-over switches on redundant networks, redundant NICs, and cables connecting redundant NICs to redundant switches.

1 All connectors have cardinality, which specifies the maximum number of
2 connections allowed per connector. For example, an IDE connector has
3 cardinality of two, one master and one slave device. See Fig. 62.

4 Principles for defining fundamental types:

- 5 ■ What is the fundamental hardware protocol?
- 6 ■ At the hardware level, what language does the device speak?
- 7 ■ Fundamental entities have exactly one owner.
- 8 ■ Connector and Connection categories must match.
- 9 ■ Duals are fail-over pairs that must be allocated as one.
- 10 ■ Entities, Connectors, or Connections can be duals.

11
12 What are the modeling elements?

- 13 ■ Entity
- 14 ■ Connector Src = Entity
- 15 ■ Connection Src = Connector, Dst = Connector

16
17 What are the fundamental categories?

- 18 ■ Entity Categories:
 - 19 ■ X86 PC: describes Software/CPU/RAM interaction. CPUs and
 - 20 RAM are values.
 - 21 ■ EFI PC: describes Software/CPU/RAM interaction. CPUs and
 - 22 RAM are values.
 - 23 ■ Network Device. Speaks IP + SNMP. Product identifier is a
 - 24 value.
 - 25 ■ Disk. Sends and receives sectors.

- Physical Container.
- Connector/Connection Categories:
 - Ethernet. Bandwidth is value.
 - ATA. Bandwidth and format are values.
 - SCSI. Bandwidth and format are values.
 - Power.
 - Physical (Containment).
 - Others: FibreChannel, Serial, Parallel, USB, FireWire, 802.11, Infiniband.

Initial Physical Configuration – See Fig. 63.

Detailed Example – See Figs. 64 and 65.

Location-Based Device Identifiers

Every networked device has a unique identifier of its location in network.

At each level, value = port number on parent switch.

Terminated levels have a termination value, “#”.

The termination value, “#”, is larger than all port numbers.

For example, see Fig. 66.

Calculating the Path between Two Devices

Consider two devices (2,0,1) and (2,1,#)

For each device, compute terminated prefixes:

1 (2,0,1) → (#,#,#), (2,#,#), (2,0,#)

2
3 (2,1,#) → (#,#,#), (2,#,#)

4
5 Most specific common terminated prefix is common parent:

6 (2,#,#)

7
8 Remaining terminated prefixes are name intermediate switches:

9
10 (2,0,1) → (2,0,#)

11 (2,1,#) → none.

12
13 Final Path:

14
15 (2,0,1) to (2,1,#) → (2,0,#), (2,#,#) = two switch hops = three wire
16 hops.

17
18 Also trivial to find closest peers to a device:

19 (2,0,1) → (2,0,?)

20 (2,1,#) → (2,?,#)

21
22 See Fig. 67.

Modeling Resource Requests

The BIG Resource Manager models the BIG machine as a graph of nodes (resources) and edges (relationships). Both nodes and edges may be annotated with attributes (name-value pairs).

One of the most common types of query against the resource graph is sub-graph isomorphism. The client creates a request graph and asks the Resource Manager to find a sub-graph within the hardware resource graph with the same shape and properties. The Resource Manager finds a match and returns a fully annotated reply graph.

As part of sub-graph isomorphism the Resource Manager **MUST NOT** fold or combine graph nodes. That is, if the request graph contains two PC nodes, the reply graph must contain two PC unique nodes.

Request graphs may include search parameters, such as find a PC node or find a PC node with at least 256MB of RAM. Reply graphs contain specific ids of each of the matching elements (both nodes and edges).

In the base case, request graphs are read-only queries. However a common optimization allows for read-write operations in the form of resource allocation. When drawn on paper, write operations are labeled with brackets.

Fig. 68 is a request graph to allocate a PC and an attached disk connected through a storage transport such as IDE or SCSI. Note that nodes are represented as round-edged rectangles and edges are represented as dark lines with overlaid rectangles where attributes are specified. The successful allocation might result in the reply graph of Fig. 69.

Driving Scenario

Joe's Flower Shop makes the resource request shown in Fig. 70. MSN insures that Joe gets at least a 500MHz PC because he has the "gold" SLA and that his PCs are attached to Switch5 to maintain locality. With the addition shown in Fig. 71, Exodus guarantee that MSN always gets machines in Rack17 and also gets small disks because they have a "2nd" class storage SLA. See Fig. 72.

Implementation Ideas

```
class Graph;

class Client
{
    private IResourceMediator mediators[];
    private Object mediatorStates[];
}

interface IResourceMediator
{
    public void MediateRequest(ref Graph graph, ref Object state);
    public void MediateReply(ref Graph graph, ref Object state);
}

class ResourceManager
{
    public Graph Allocate(Graph request, Client client)
    {
        for (int n = 0; n < client.mediators.Length; n++)
        {
            client.mediators[n].MediateRequest(ref request,
                ref client.mediatorStates[n]);
        }

        Graph reply = PrimitveAllocate(request);
    }
}
```

```

1         for (int n = client.mediators.Length - 1; n >= 0; n--)
2         {
3             client.mediators[n].MediateReply(ref reply,
4                 ref client.mediatorStates[n]);
5         }
6         return reply;
7     }
8     private Graph PrimitiveAllocate(Graph request);
9 }

```

Basic Resource Allocation Scenarios

This section list a number of scenarios. Included with each scenario is the corresponding request graph. Nodes that will be allocated as a result of the query transaction are labeled with “[Allocate]”. Nodes that will not be allocated and that must be unallocated for the search to match are labeled with “[Free]”. Nodes without a bracketed label are not allocated, instead they provide context for the rest of the request graph.

PC

Akamai needs to allocate a server in the Digix data center with at least a 1GHz CPU, 512MB of RAM, and 100GB of local disk storage. See Fig. 73.

VLANs

1 MSN Instance Messaging has decided to implement a DMZ containing its
2 front-ends. In order to do so, it needs 2 VLANs with coverage of its front-ends.

3 See Fig. 74.

4 5 Public IP Address or DNS Name

6
7 Joe's Web Service needs to make itself visible to the outside world. He
8 needs to allocate a DNS entry and a routable IP Address. See Fig. 75.

9 10 Load Balancing Groups

11
12 Joe's Web Service has grown too large for a single PC. He needs to
13 allocate a load balancing group and another PC. He then needs to place both PCs
14 behind the load balanced group's virtual IP address. See Fig. 76.

15 16 Path

17
18 Hotmail needs to allocate an 80Mbps path to transfer email accounts from
19 one UStore to another. Hotmail may specify latency and QOS requirements on the
20 path also. See Fig. 77.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

Specific Storage

Hotmail wants to create a new UStore. It wants a Raid 1 box with 100GB spread over at least 4 sets of unshared heads rotating at 10,000 RPM or better. See Fig. 78.

Cluster (Quorum) Storage

Hotmail wants to allocate a pair of machines with a shared disk for a fail-over cluster. It wants a Raid 1 box with 100GB spread over at least 4 sets of unshared heads rotating at 10,000 RPM or better. See Fig. 79.

Shared Storage

Joe's Web Service needs 50GB of common storage usable by multiple machines to hold rollback images of service specific configuration. The storage is available to 0 to N machines. See Fig. 80.

Allocation Placement Scenarios

Proximal Machine Allocation

Hotmail needs to allocate a new front-end. It wants to find a machine on the same switch as its other front-ends with sufficient bandwidth to the back-end cluster. See Fig. 81.

Distant Machine Allocation

The Expedia customer profile database needs to another machine for SQL replication. It wants a machine that is located in a part of the data center covered by a different battery backup unit. See Fig. 82. Or possibly the example of Fig. 83.

Latency Driven Allocation

The Hotmail back-end needs to allocate a machine for cluster coordination. The machine must be within 5ms latency of the machines already in the cluster, but bandwidth is low. Alternatively this could be represented by needing the machine to be within 1 network hop. See Fig. 84.

Seeding a Compound Component

1 Hotmail is about to create a new email unit. The unit should be allocated in
2 a single-hop cluster with room to grow to at least 500 PCs, although Hotmail may
3 only initially allocate a few dozen machines. See Fig. 85.

4 5 Batch Allocation

6
7 MSN Search decides to add the ability to search MP3s for based on small
8 music samples. It wants to allocate a block of 400 PCs, 3 load balancers, and
9 20TB of storage. It wants an all-or-nothing allocation. See Fig. 86.

10 11 Revocation Scenarios

12 13 Recovery

14 Joe's Web Service has stopped paying the IDC. The IDC needs to recover
15 all of the resources allocated to Joe's Web Service and return them to the pool of
16 available resources.

17 18 Hardware Lifetime Revocation

19
20 One of Expedia's front-ends is a PC which has reached the end of its life
21 cycle. The triggered by the IDC's operation logic, the Resource Manager notifies
22 Expedia that it has 72 hours until the machine is returned to the IDC.

Controlled Revocation

Hotmail allocated 20 short-term machines for a massive reshuffling of its UStores. In accordance with its SLA, the IDC has now asking for one machine to be returned. Hotmail can either return one of the twenty or another equivalent machine.

BIG Vision - Enable:

- development of distributed, scalable and highly available services using Visual Studio and reusable building blocks like SQL, IIS, ...
- deployment across a set of abstracted hardware and software resources that are automatically allocated, purposed and configured
- lower cost of ownership through automation by codifying operational best practices to control service availability and growth
- procurement of standardized data center hardware that leverages commodity economics

BIG Services Platform Architecture – See Fig. 87.

BIG Computer – Hardware Reference Platform

Reduces the cost of design, test and operations:

- Limits number of hardware devices to support
- Constrains the network topology
- Enables automation of network configuration

Eliminates customer concerns about BIG technology deployment requirements

- PXE, DHCP, DNS, VLANs

IP Gateway

- Mediates IP traffic between the external network and the internal network
- Network Address Translation (NAT), firewall, load balancing

Internal Network

- IP addrs and VLANs are managed exclusively by BIG
- VLANs are automatically configured

Hardware Building Blocks

- Combinations of commodity servers, network switches, and disks

See Fig. 88.

Fig. 89 illustrates examples of current products that can be inside a BIG computer.

Resource Management Features

- Dynamic discovery of server, storage or network hardware resources.
- Highly available database containing (physical and logical) resources.
- Runtime API that supports enumeration, querying and updating of resources.
- Logical resource driver model and API for binding resource drivers to physical hardware devices.
- Programmatic allocation and de-allocation of server resources.
- Automatic configuration and management of network resources such as VLANs and load balancing groups.

- Dynamic configuration and management of block and file-based storage resources.
- Failure detection monitoring and notification.

Resource Management Components

- Resource Managers are responsible for allocation of hardware and software resources inside the BIG Computer
 - Resource managers register with the BIG runtime
 - Resource managers are essentially factories for a given resource type
- Hardware Resource Manager
 - Base level factory responsible for allocating hardware instances
- Network Resource Manager
 - Responsible for allocating VLANs, Load balancing groups, IP addresses, ...
- Storage Resource Manager
 - Manages storage resources such as disks and files
- PC Resource Manager
 - Allocates target servers and deploys OS using iBIG services
- Software Resource Managers
 - Allocates and configures IIS vroots, SQL databases, ASP .NET, ...

Fig. 90 illustrates various resource management components.

Hardware Resource Discovery and Management

Properties: Power, Network, Storage, Processor, Memory, Location

1 Hardware inside BIG Computer is automatically discovered. Resource drivers are
2 bound to hardware devices and expose logical resources to Hardware Resource
3 Manager (HRM). HRM translates a logical resource allocation request to a
4 physical resource binding. See Figs. 63, 64, and 65.

5 6 Network Resource Management within the BIG Computer

7 BIG Computer defines an abstraction layer for network resources.

8 Network Resource Manager: allocates network resources and programs the
9 network switches and load balancers inside the BIG Computer, and interfaces with
10 the network resource drivers.

11 VLANs provide isolation and partition the network inside the BIG Computer.

12 Network resource examples: VLANs, Load Balancing Groups, Network Filters,
13 IP addresses, DNS names.

14 15 BIG Storage Resource Management Requirements

- 16 ● Global view of storage connected to the BIG Computer that encompasses
17 file and block-based storage resources.
- 18 ● Virtualization of the storage interconnect fabric.
- 19 ● Framework for creating and managing higher level storage abstractions
20 such as LUNs, volumes, arrays, etc.
- 21 ● A driver/provider model to allow existing and new storage devices to plug
22 into a BIG Computer.
- 23 ● Interoperability with SAN systems.

Infrastructure Services (Automated Deployment Services (ADS)) - Features

- Base Deployment Services

- Basic Network Boot Service (PXE) and Image Builder Service
- Pre-boot OS environment (BMonitor)
- Virtual floppy delivered over network for legacy tools support

- Image Deployment and Management

- Tools for creating, editing and deleting images
- Deployment of images to systems running pre-OS

- Multiple Device Management (MDM)

- Scripts for common tasks
- Task sequencing to coordinate multiple steps and processes for deployment
- Full programmatic interface (WMI)

- Ships 60 days from .NET Server RTM

- Supports Windows 2000 and .NET Server targets

Fig. 92 illustrates an example ADS Architecture.

Fig. 93 illustrates an example ADS Remote Boot and Imaging system.

Service Definition Model (SDM)

- The programmatic description of the entire service

- Declarative definition of the service
- Defines the overall service structure of the service in a scale-invariant manner
- Provides a framework for deployment, management, and operations

- 1 ■ Component-based model captures in a modular fashion the elements
- 2 of a service
- 3 ● SDML is the declarative language for defining Service Definition Models
- 4 ■ Components, ports and wires
- 5 ■ Type, member and instance space
- 6 ■ Supports composition and encapsulation

8 SDM: Components, Ports and Wires

- 9 ● Components are units of implementation, deployment and operations
- 10 ■ For example, dedicated server running .NET Server, IIS virtual web
- 11 site, SQL database
- 12 ■ Expose functionality through ports and communicate through wires
- 13 ■ Compound components created by composition
- 14 ● Ports are names (service access points) with an associated type (protocol)
- 15 ■ BIG does not mandate what protocols to use for communication
- 16 ■ Protocols capture the information required for establishing
- 17 communication
- 18 ● Wires are the permissible bindings between ports
- 19 ■ Wires declare a topological relationship between ports

20 See Fig. 94.

21 Fig. 95 illustrates an SDML example: MyService.sdml. Fig. 28 is also related to
22 this SDML example.

1 Service Deployment Unit (SDU) – Encapsulates all the pieces that make up a
2 service, including: SDM model for the application/service, CLR assemblies for
3 component implementations, and MSI, ASP.NET, SQL scripts, Static content, etc.
4 See Fig. 96.

6 SDM Runtime

- 7 ● SDM Runtime is responsible for tracking SDM models and instances
 - 8 ■ Implemented as a Web Service hosted by IIS
 - 9 ■ Can be partitioned for scalability
- 10 ● Runtime API exposes SOAP endpoints
 - 11 ■ Communication with runtime is done through a runtime library
- 12 ● Highly available SDM Store (using Yukon's redundant database
13 technology)
 - 14 ■ Two SQL servers and a witness server

15 See Fig. 27.

17 Example: Component Instantiation
18 using Microsoft.SDM;

```
19 public class MyService:
20     SDMComponent
21 {
22     public OnCreate(...) {
23
24         fe1 = CreateInstance("fe", "");
25         be1 = CreateInstance("be", "");
26
27         w1 = CreateWireInstance("tds");
28         w1.Members.Add(fe1.Ports["catalog"]);
29         w1.Members.Add(be1.Ports["sql"]);
```

```

1      }
2    }
3
4    myservice.cs is C# code that uses the SDM API.
5
6    componenttype MyService
7    {
8        component MyFrontEnd fe;
9        component MyBackEnd be;
10       port http = fe.http;
11       wire TDS tds {
12           fe.catalog;
13           be.sql;
14       }
15       implementation "MyService, MyCLRApp"
16   }

```

See Fig. 35.

Example of Dynamic Binding using SDM Runtime APIs (See Fig. 97)

1. be[1] declares that sql[1] port is ready and registers its port connection information with the SDM Runtime using DeclarePort()
2. fe[1] initializes and asks the SDM Runtime for peer information for catalog[1] port and receives information about sql[1] port using GetPeerPort()
3. fe[1] then connects to be[1] using the port connection information provided dynamically by the SDM Runtime

Service Definition Model (SDM) Workgroup

- SDM Workgroup is comprised of 5 teams
 - Indigo
 - Whitehorse
 - Fusion

- 1 ■ Management
- 2 ■ BIG
- 3 ● Charter was to define a class-level application schema for distributed and/or
- 4 heterogeneous applications
- 5 ■ Describes applications using components, ports and wires
- 6 ■ Includes deployment, configuration and management information
- 7 ● SDM is an exoskeleton that references Fusion and Management (and
- 8 potentially other) schemas
- 9 ■ Fusion assemblies are referenced for deployment (where applicable)
- 10 ■ MBU Settings and Instrumentation schema are referenced and
- 11 specified for configuration and monitoring

13 SDM Schema (simplified)

```

14  <sdm>
15    <identity />           // identifies the group of definitions
16    <porttypes />         // descriptions of ports
17    <wiretypes />         // descriptions of topologies
18    <componenttypes>      // set of components defined in this library
19      <componenttype>
20        <ports />         // communications capabilities
21        <settings />      // configuration settings for component
22        <instrumentation /> // monitoring schema
23        <deployment />    // installer type, installer info, (e.g., Fusion)
24        <components />    // subcomponents for composition
25        <wires />         // defines relationships between ports
26      </componentType>
27    </componenttypes>
28  </sdm>

```

SDM and Fusion – See Fig. 98.

- Local settings with default values are specified in the Fusion Manifest (or other local install technology).
- Settings in an SDM are processed by Ops Logic and the BIG runtime.
 - Example: “number of users” would be used to determine the initial scale-out condition of the application

SDM and Deployment – See Fig. 99.

Describing the structure of an application in a scale-invariant manner requires a similar scale-invariant description of the application host environment (i.e., data center) to enable design-time validation of deployment requirements and constraints.

- Microsoft and customers expend lots of energy drawing elaborate descriptions of their data center environments and writing very large documents to explain the drawings.
- These drawings and documents merge many layers of information from physical machine names to IP addresses to VLANs to server roles into one comprehensive view that is often confusing.

Fig. 100 illustrates an example system architecture.

Fig. 101 illustrates an example of various deployment layers.

1 Operations Logic is the “Business Logic” of Operations

2 Operations Logic is CLR code that captures repeatable patterns encoded as
3 reusable best practices

- 4 ■ Not specific to a service or operating environment
- 5 ■ Can be developed, tested and shipped
- 6 ■ Reduces the need for manual procedures that require people to
7 execute them

8 OpsLogic is responsible for the overall operation of a service

- 9 ■ Starting up a service
- 10 ■ Service growth and shrinkage
- 11 ■ Upgrades and updates
- 12 ■ Fault detection and recovery
- 13 ■ Database partitioning

14 OpsLogic will be implemented using MS middle-tier technologies

- 15 ■ ASP.NET web services hosted on IIS
- 16 ■ DTC for transaction coordination
- 17 ■ SQL server for storage
- 18 ■ WMI for monitoring and management
- 19 ■ MSMQ for messaging

20
21 Repeatable Upgrade Patterns → Operations Logic

- 22 ● Upgrade is an example of the type of reusable Operations Logic template
23 we want to ship with BIG
- 24 ● In-place Upgrade Pattern

25

- 1 ■ Cost of moving data is high, code instantiation cost is low, or no
- 2 spare resources
- 3 ■ Takes component out of service, runs update, put it back in service
- 4 ● Side-by-side Upgrade Pattern
- 5 ■ Cost of moving data is low, code instantiation cost is high, have
- 6 spare resources
- 7 ■ Create new component; Take old component out of service; Migrate
- 8 data to new component; Put new component into service
- 9 ● Replacement Upgrade Pattern
- 10 ■ No data migration
- 11 ■ Add new components; remove old ones; coordinate to maintain
- 12 service availability
- 13 ● Rolling Upgrade is an example of higher-level operations logic that can
- 14 reuse the codified upgrade patterns
- 15 ■ Operations logic can be tested and the framework supports rollback
- 16 ■ Removes human error from execution by letting software perform
- 17 the steps

18

19 Operations Logic, BIG and the Microsoft Programming Model - See Fig. 102.

20 The Internet transforms enterprise applications – Increased exposure has

21 resulted in increased costs. See Fig. 103. New architecture has led to an increase

22 in costs driven by HW, people and a decrease in agility due to complexity.

23 See Fig. 104. Moore's Law is spreading across the DC – dramatic increase in disk

24 density, NW throughput and processing power.

25

1 Service delivery is people intensive – human involvement impacts security,
2 reliability, flexibility and cost. See Fig. 105.

3 This is a lifecycle problem – customer pain spans develop, deploy, and
4 operate phases. See Fig. 106. Applications are not developed with: scale in
5 mind-tied to HW configurations, manageability in mind, operations in mind- what
6 are the requirements in my data center? Test – “Thrown over the wall”.

7 Developer desktop → test configuration? How does this map to my production
8 environment. Deployment challenges: Which servers do I use? What is the right
9 topology? Have I checked with the server, storage and networking teams? How
10 much future demand do I need to anticipate? Operational challenges: What do I
11 do with all of these alerts? How will that failing NIC affect my application? Why
12 is the performance of my service degrading? I wish I could clone my email
13 admin.

14
15 Addressing the service delivery challenges – core tenants of a viable solution for
16 customers.

17 Independent value at each step of the lifecycle

18 Develop, Deploy, Operate

19 Unifying architecture for the entire lifecycle

20 Improved coordination and feedback between steps

21 Enable mapping to changing business needs

22 Mapping can only be done once you have agility

23 Built on lowest TCO platform

24 Effectively leverage industry standard hardware through scale out
25

1 Project Summit – A revolutionary service delivery architecture. See Fig.
2 106. Develop services that: are instrumented and manageable, include
3 deployment requirements, encapsulate operations knowledge, and leverage
4 standard building blocks. Easily deploy services: rapid provisioning, DC resource
5 virtualization, self contained, one-click deploy, consistently from test to
6 production, and independent of scale. Simplified operation: aggregated
7 administration, monitoring and change management, manage services not
8 applications, true automation via context, rich service-centric management
9 console.

10 Map business needs to IT systems. Capture IT operational knowledge in
11 the tools.

12 Project Summit – a comprehensive new architecture and an industry wide
13 initiative. See Fig. 107.

14
15 Concept → Architecture → Product

16 A long term, customer and partner-driven effort.

17 A major investment beginning in 1999

18 Began with deep research into operational needs of large MS internet
19 properties

20 Validated initial finding across broad customer base

21 Prototype from the product group in late 2000

22 Strong set of joint development partners

23 Large enterprise and service provider customers involved in product
24 definition
25

IHV and ISV partners consulted to help define functionality exposed via
APIs

Initial product shipping with Windows Server 2003

Customers convert complex systems into simple diagrams. See Fig. 108.

Who is involved in delivering your IT Service? – Humans are an integral part of
the system.

Application architect – designs the service.

Network architect – configures the network.

Storage architect – configures remote storage.

Application operator – maintains the service.

Network operator – maintains the network.

Storage operator – maintains remote storage.

Server operator – maintains the servers.

Problems with this model: many human interactions, no common language,
blurring of domain knowledge.

Details of the solution:

Service Definition Model

Resource Virtualization

Operational Automation

Management APIs and Solutions

Driving an Industry wide initiative

1 The Service Definition Model (SDM) – capturing the complete service.

2 Comprehensive description of a service

3 Application components and instrumentation

4 Service topology

5 Underlying resources (server, storage, network)

6 Relevant to developers and operators

7 Layers and separates responsibility

8 Provides a consistent frame of reference

9 Exposed in Visual Studio for developers

10 A living model at run time for operators

11 Logically consistent ind. of allocated resources

12 Tracks resources in real time

13 The single authority on service composition

14 Provides context for true automation

15
16 **SDM Terminology**

17 **Components – the building blocks of services.**

18 Logical construct

19 Scale invariant

20 One component may have multiple instances

21 Simple or compound

22 Single logical entities (database, web service, file partition)

23 Combined logical entities (HA database, email, etc...)

24 Include a deployment manifest specific to the component

25 DB component includes the database schema

1 Web service component includes URL directory, content, code
2 Interconnected with ports and wires
3 Ports - service access point
4 Wires – communication relationship between ports
5
6 SDM provides the means for abstraction and encapsulation. See Fig. 110.
7 Enables reuse
8 Structures complexity
9
10 Mapping people to the SDM – provides a consistent frame of reference. See Fig.
11 111.
12
13 Developing an SDM application – a new Visual Studio design surface. See Fig.
14 112. Legacy apps, New apps.
15
16 An SDM service in the data center – comprehensive description with living model
17 tracking resources. See Fig. 113.
18
19 What is a Summit Computer?
20 An agile pool of virtualized hardware resources
21 Servers, Storage, Network Devices, Managed Fabrics.
22 Few dozen to few thousand servers.
23 Assembled from existing HW or ordered as one SKU from OEM.
24 A single managed entity
25

1 Summit provisions and manages all HW resources w/in Summit
2 Computer.

3 Summit owns complete configuration of internal network fabrics.

4 A bounded domain of control

5 Standardized topology bounds complexity of build, test, and
6 operations.

7 Ownership unchanged for resources outside the Summit Computer.

8 A catalyst for software innovation

9 Q: What data center environments should I target for my server
10 application?

11 A: The Summit Computer.

12 Just like Win3 let ISVs forget about details of printers and
13 graphics cards.

14 A catalyst for hardware innovation

15 Microsoft engaged with major hardware vendors to define a
16 reference platform.

17 First specs & innovations to appear at WinHEC (May 2003).

18 Summit provides SW environment for aggregation innovations:

19 Dense blades, Smart racks, etc.

20 Summit enables simplification of hardware, for example allows:

21 Drop KVM from servers and human interfaces from network
22 devices.

23
24 Fig. 114 illustrates example resource managers.
25

1 Resource Virtualization – the bridge between the SDM and component instances.

2 Responsible for sharing, allocating, and recovery. See Fig. 115.

3
4 Server Resource Virtualization – Automated Deployment Services (ADS) in
5 Windows Server 2003.

6 Complete infrastructure for rapidly purposing and re-purposing Windows
7 Servers

8 Imaging tools to capture and edit both Windows 2000 and Windows
9 Server 2003 images

10 Secure, remote deployment framework enabling zero touch server
11 builds from bare metal

12 A framework for mass server administration

13 Secure, reliable, script execution infrastructure

14 Programmatic model of your Windows datacenter

15 A persistent log of all administrative activities

16 Graphical and Programmatic interfaces

17 Simple MMC UI for GUI-based operation

18 Full functionality exposed through command line tools and/or WMI
19 layer

20
21 Key benefits of ADS

22 1. Lower the TCO associated with bare metal server builds and script-based
23 administration

24 Enable zero-touch server builds from bare metal

25 Secure script based administration of 1000 servers as easily as 1

1 2. Improve the consistency, security and scalability of your Windows Server
2 datacenter

3 Encode operational best practices and eliminate human error

4 Maintain a persistent store of all administrative activities

5 Centrally perform secure, script-based administration of your entire
6 Windows datacenter

7 Rapidly change server role in response to changes in workload
8 requirements

9 3. Leverage your existing server administration investments

10 Extend and enhance your existing script-based automation
11 methodologies

12
13 Operational Automation – core tenets of automation

14 Flexible framework to enable capture and re-use of operational best practices

15 Operations Logic

16 Rich context within which to automate

17 Events are contextualized by the SDM to enable true automation of
18 systems management

19 “Which application will be effected by the NIC that dies on the 5th
20 DL380 in rack 22?”

21 Transact-able

22 Compensation based model allows rollback and un-do

23
24 Operations Logic – a framework for developer and operator automation.

25 What is operations logic?

1 Encoded operations processes that are long-lived, highly-available and
2 durable

3 Leverages the SDM for context and control of Summit computer
4 resources

5 Enables operators to vary the level of automation in a system

6 Benefits for the developer

7 Allows the developer to capture how the system should respond to and
8 resolve application events and messages (such as return codes)

9 Enables Microsoft and the ISV community to provide standard,
10 predefined operational processes that the developer can use or extend

11 Deploy, upgrade, scale-out and remove resources

12 Benefits for the ITPro or Operator

13 Enables easy re-use of proven operational best practices for the data
14 center

15
16 Operational Automation – programming operations logic. See Fig. 116.

17 How the SDM interacts with operations logic:

18 Events are annotated to indicate instance and component information

19 The monitoring subsystem does time-based event correlation

20 Alerts are a roll-up of events

21 Greater semantic meaning

22 Commandlets are,

23 The set of management commands exposed by a component

24 Are self-describing

25 Can be used directly within a shell

1 Can have a GUI forms representation

2 Can provide a “man-page” for use by operators

3 See Fig. 117.

4
5 Operational Automation – transact-able.

6 Transactions are essential to support fault-tolerant operations

7 Example: Adding a web service

8 Powerful extension to ad-hoc shell scripts

9 All forms of Operations Logic function under the auspices of a transaction
10 model

11 Compensation-based

12 Durable

13 Using orchestration, transactions can span multiple machines

14
15 Management APIs and solutions – leveraging the richness of the SDM.

16 Visualization occurs through the SDM

17 3rd party consoles can pull information directly from the SDM or
18 leverage platform know-how of the Microsoft management solutions
19 Microsoft will build an SDM-based management console for the data
20 center

21 Customer could create custom consoles via the SDM

22 See Fig. 118.

23
24 Industry wide initiatives – unleashing IHV, ISV, SI innovation.

25 IHV HW reference platform specification

1 Working closely with major OEMs and switch manufacturers

2 Targeting release at WinHEC (May '03)

3 Drive new compelling features into future HW offerings

4 Engage key 3rd party ISVs

5 Create application components for Visual Studio

6 Resource managers for their applications within the SDM

7 Mgmt ISVs to create SDM based management consoles

8 Work with SIs as both customers and partners

9 Customer

10 Dramatically lower their operational costs

11 Partner

12 Create innovative new service offerings on this platform

13 Capitalize on operations expertise → develop Operations Logic

14
15 Major Customer Benefits: provide choice and create the most economical,
16 manageable platform for the data center.

17
18 Industry wide initiatives – extending the richness of the SDM to heterogeneous
19 environments. Developing heterogeneous SDM applications using Visual Studio
20 (enables development of SDM applications for Windows) or 3rd party tools (enable
21 development of SDM applications for other platforms).

22
23 Conclusion

24 Although the invention has been described in language specific to structural
25 features and/or methodological acts, it is to be understood that the invention

1 defined in the exemplary appended claims is not limited to the specific features or
2 acts described. Rather, the specific features and acts are disclosed as exemplary
3 forms of implementing the claimed invention. Moreover, these claims are
4 exemplary in terms of scope and subject matter. Many other combinations and
5 sub-combinations of the features described herein may later be claimed in patent
6 applications claiming priority to this application.

7 Although the description above uses language that is specific to structural
8 features and/or methodological acts, it is to be understood that the invention
9 defined in the appended claims is not limited to the specific features or acts
10 described. Rather, the specific features and acts are disclosed as exemplary forms
11 of implementing the invention.